Performance under Pressure in the NBA

Zheng Cao
Oregon State University

Joseph Price
Brigham Young University

Daniel F. Stone
Oregon State University

June 14, 2010

Abstract

We analyze the effects of pressure on performance using National Basketball Association (NBA) free throw data from the 2002-03 through 2009-10 seasons. We find strong evidence that players choke under pressure—they shoot 5-10% worse than normal in the final seconds of very close games. Choking is more likely for players who are worse overall free throw shooters, and on the second shot of a pair after the first shot is missed. In general, performance declines as pressure increases (as game time remaining decreases, and as the score margin decreases, whether the shooter's team is winning or losing). However, we find no evidence of choking when games are tied in the final 15 seconds. We also fail to find evidence of performance under pressure being affected by home status, attendance, and whether or not the game is in the playoffs.

JEL Classification Numbers: L83, D03, M54, J24

Keywords: Choking, Anxiety, Pressure, Clutch, NBA

Daniel F. Stone (corresponding author) –Department of Economics, Oregon State University, 303 Ballard Extension Hall, Corvallis, Oregon, USA, 97331, dan.stone@oregonstate.edu, (541) 737-1477 (phone). Zheng Cao—Department of Economics, Oregon State University, caoz@onid.orst.edu. Joseph Price—Department of Economics, Brigham Young University, joe_price@byu.edu, (801) 422-5296**.**

Electronic copy available at: http://ssrn.com/abstract=1624975

I. Introduction

Neoclassical economic theory predicts that individuals exert the most effort, and consequently produce their best performances, when the returns to effort are highest. Folk wisdom and common sense suggest something different: when performance matters most individuals feel psychological pressure, and as a result often make uncharacteristic mistakes (i.e. "choke"), but sometimes rise to the occasion in the clutch. The degree to which experienced professionals choke or perform well in the clutch, and the factors causing or mitigating either type of behavior, are empirical issues that are still not fully understood.[1]

In this paper we analyze performance under pressure using free throw data from the National Basketball Association (NBA). Pressure and the potential for choking arise in all sports, and many contexts outside of sports as well.[2] The basketball variable free throw percentage is a measure of performance that is especially well-suited for the empirical analysis of choking for several reasons. First, each free throw attempt is taken from the same distance and location, so the physical difficulty of each shot is constant (as opposed to basketball field goal attempts or say, golf putts, which are taken from varying locations). Second, free throws are undefended, and so observed changes in performance are not confounded by simultaneous changes in offensive and defensive player behavior (in contrast to, e.g., soccer penalty kicks). Third, free throws occur very frequently—most players shoot over 100 in each season. This allows us to control carefully for the substantial heterogeneity across players in free throw shooting ability. Fourth, since most close games involve at least a few free throws being shot in the final seconds or minute, we have access to the sample size necessary to obtain reasonably precise estimates of how performance changes as a function of time remaining in the game and score difference. Finally, shooting a free throw is a fairly easy, but still non-trivial, task for most NBA players (as opposed to, say, the trivial task of kicking an extra point in American football), so *a priori* it seems highly plausible that psychological factors may affect performance.

We find that NBA players do indeed choke, but not too badly. In fact, they are fairly immune to factors that we hypothesized may cause anxiety throughout games. In the first three

[1] For example, in a recent Newsweek blog post on the psychology of missed field goals the author wrote, "All of this raises the question: could the preceding outbreak of failed field-goal attempts have precipitated Kaeding's spectacular meltdown? Did Kaeding fall prey to a shanking epidemic? According to Dr. John F. Murray, a Palm Beach-based sports psychologist, it's a plausible theory, although impossible to prove," (Yarett, 2010).

[2] See Dohmen (2008) and Ariely (2010) for discussions of the relevance of choking in labor market contexts outside sports. One context particularly worth noting is finance, and perhaps day-trading in particular, since, like basketball players, traders sometimes have to make split-second decisions and the degree of pressure can be very high.

Electronic copy available at: http://ssrn.com/abstract=1624975

quarters of games, being the home team has a statistically significant but very small (less than 0.5%) positive effect, playoff status has no effect, and attendance has a significant, but also small effect (free throw percentage decreases by just over 1% when attendance increases by 10,000) for both home and away players. The score difference and time of shot have small (1% or less) and generally insignificant effects.

When the shot is taken at the very end of the game, and thus has the potential to directly change the game's outcome, the story is different. We find that expected free throw percentage declines by around 4% when the shooter's team is down by one or two points and the shot is taken in a game's final minute. The magnitude of the decline increases to 6.3% and 8.8% in the last 15 seconds, when down two and one points, respectively. There are also significant choking effects for players whose teams are up by one or two points, but they are substantially stronger for the second shot, when pressure is likely higher.[3] Surprisingly, given these other results, we find that performance does not decline at all in the last 15 seconds when the score is tied.

We find that choking effects are considerably larger for players who are worse free throw shooters in general. Since it is natural to think that these shooters are less confident in their abilities, we interpret this result to support the theory that self-confidence moderates, and potentially eliminates anxiety and choking. This result may also simply be due to the magnitude of choking being an increasing function of the variance of the outcome variable. We also find that choking declines when the previous free throw was made and increases when the previous shot was missed, which provides further support for the importance of confidence, since making the first free throw likely improves confidence.[4] We find that the effects of other variables that may be thought to affect the degree of choking in games' final seconds (experience, home and playoff status, attendance) are minimal and mostly insignificant.

Our results stand in sharp contrast to the findings of Dohmen (2008), the only published sports economics paper we are aware of that empirically analyzes choking. We discuss the relation between our papers, and other related literature in Section II. We discuss our data and empirical methods in Section III, results in Section IV, and Section V provides additional

[3] Most free throws are shot in pairs.

[4] Arkes (2010) showed that free throw shooters perform better after making their previous free throw, as compared to after a missed free throw, which he interpreted as evidence of the hot hand. We find similar results, and find that the effect is exacerbated considerably in the final seconds of close games. It is also worth noting (estimates not reported) that players perform substantially (3-4%) better on the second shot than the first even when the first was missed, indicating that even seasoned professionals benefit from a warm-up shot.

interpretation of results, including speculation regarding the finding that choking does not occur when the score is tied, and concluding remarks.

II. Theory and Related Literature

The topic of choking under pressure has received substantial attention from the psychology literature. See Hilla et al. (2010) for a recent review. The authors say that most of the literature defines choking (at least implicitly) as "any inferior performance under pressure."[5] We attempt to use this basic definition, which raises two questions: 1) What causes pressure? 2) Inferior performance as compared to what? Regarding the former, it is natural to think that pressure increases as the probability that the performance will affect which team wins the game increases, which occurs as both the score difference and time remaining in the game decrease. Hilla et al. (2010) also discuss what they refer to as possible moderators of choking, including presence of an audience, whether that audience is supportive or not, self-confidence and skill level.[6] The answer to the second question is less straightforward, since it is not clear that pressure and performance have even a monotone relationship. Hilla et al. (2010) discuss the upside-down U curve theory of this relation: that increasing pressure first causes performance to improve, then decline. We take this into consideration in our empirical analysis.

Many of the papers cited by Hilla et al. (2010) use experimental data from controlled settings; one advantage of our paper is that we use non-experimental data.[7] The size of our dataset (discussed further below) also provides us with greater statistical power than that of many experimental studies. One especially closely related psychology paper we should discuss further is Worthy et al. (2009), since it also analyzes NBA free throw performance and uses a fairly large dataset. Their results are fairly similar to some of ours, in that we both find performance

[5]The authors provide a detailed discussion of alternative theories for exactly why, and how, anxiety may arise and worsen performance in pressure situations. Two primary classes of theories are distraction (that performers are distracted from their task when pressure is high) and self-focus (that performers become self-conscious under pressure). Distinguishing between these theories is beyond the scope of this paper; we refer the reader to their paper for more detail on those topics.

[6] See Wallace et al. (2006) for a review of the literature on the effects of performing in front of supportive and hostile audiences. They discuss the theory of how supportive audiences should improve performance for effort-based tasks, but increase pressure and worsen performance for skill-based tasks. Evidence of supportive audiences causing choking has been found in multiple contexts, such as golf and baseball. Free throw performance is likely more affected by skill than effort, however, basketball audiences can distract players from the opposing team, especially free throw shooters, more than audiences in other sports (such as golf), so the net theoretical effect of home status on performance in basketball is ambiguous.

[7]See, e.g., Wang et al. (2004) for a study of choking that uses free throw data obtained from a controlled setting. See Levitt and List (2007) for a discussion of the advantages of data "from the field".

worsens when the shooter's team is down by one or two points, or up one point with less than one minute remaining, but not when tied. However, there are numerous differences between our papers. We use a sample from eight seasons, while their sample only comes from one season, and we use a different empirical strategy. We also examine several issues outside of their scope, including pressure effects throughout the game, how the effects change as time remaining decreases during the last minute and depending on whether the shot is the first or last of the set, and how the effects vary for players with different characteristics.

There is a limited but growing economics literature on choking; Ariely et al. (2009) analyze experiments showing that individuals perform worse when stakes are higher, and Bannier and Feess (2010) and Rauh and Seccia (2010) study theoretical models that incorporate anxiety and choking. Dohmen (2008) is the only published paper we are aware of that uses non-experimental data to focus on the choking phenomenon.[8] His paper uses penalty kick data from the German Premier football league over 40 seasons. He finds that while being the home team is associated with an approximately 2% decline in conversion percentage, neither the score difference at the time of the shot, nor attendance, have significant effects on performance. Our results are strikingly different: we find that being the home team has very little effect, while score can have an effect larger than 2%. There are several differences between our settings that may account for the differences in our results.[9] First, penalty kicks involve offensive-defensive player interactions, in contrast to free throws. Thus, if the defensive player's behavior changes in pressure situations, this could confound the results. Dohmen (2008) accounts for this by focusing on the percentage of penalty shots that completely miss the goal (and so the conversions are not prevented directly by the goalkeeper), but this issue cannot be completely eliminated. If kickers know that choking occurs, and keepers are prone to it as well, this could reduce kickers' anxiety and reduce observed choking. Second, while obtaining the opportunity to shoot a penalty kick or free throw is endogenous in both sports, this problem may be more severe in soccer. Dohmen (2008) says, "An endogeneity problem arises if players retire from penalty kicking after having failed on a penalty kick," which does not occur in basketball (all players sometimes shoot free throws and basketball coaches only choose the shooter for technical fouls,

[8] Ariely (2010) refers to research on clutch performance using NBA data currently in progress. There are also sabermetric studies on clutch performance, but they focus on the question of whether individual players are consistently good clutch performers; see, e.g., Neyer (1999).

[9] We should also note that our settings are comparable in that the overall success rates for penalty kicks and free throws are very similar (both around 75%), so the tasks are of similar difficulty.

which are infrequent). If kickers who are better on average or better under pressure are given more opportunities to kick under pressure, this would dilute estimates of choking, especially since Dohmen's study does not control for player effects. Also, penalty kicks occur much less frequently than free throws. As a result, the number of penalty kick observations that occur at the very end of games is very small, and we find that choking effects are much stronger in those situations. On the other hand, each penalty kick is much more important than most individual free throws, since scoring occurs much less frequently in soccer than basketball. This should cause the importance of time remaining at the time of the kick, with respect to kicker performance, to be mitigated. Finally, the home crowds at soccer games may be more intense than those of basketball games, which may account for the differences in estimated home effects.

3. Data and Empirical Strategy

We use play-by-play data from the 2002-03 through 2009-10 seasons obtained from ESPN.com. We use a play-level, rather than game-level, dataset because it allows us to observe the time in the game at which the various shots were taken, and the score difference at the time of the shots. We convert the play-by-play data to a free throw attempt-level dataset with nearly 500,000 observations; however, the vast majority are shots taken throughout the game that are unlikely to directly cause the shooter's team to win or lose, and so pressure likely does not vary substantially across most of the observations. Summary statistics are presented in Table 1 and graphically in Figure 1 (all variables referred to are defined below).[10] We have over 300 and over 500 observations with the shooter's team losing by less than five points with 16-30 and 0-15 seconds remaining in the fourth quarter or overtime, respectively.[11] The table and figure also provide a preview of our econometric results, as they both show performance clearly declines as the score margin and game time remaining decrease in the final minute of games. Both indicate that choking is greater when teams are losing rather than winning (by the same margin), and Figure 1 indicates a lack of choking in the second to last minute of games, and that choking does not occur when the score is tied in the final 15 seconds. A more subtle thing to note from Figure

[10] Attendance data are only available for the 2002-03 through 2007-08 seasons. We began this project during the 2009-10 season and added in all observations for the regular season as we proceeded, but exclude all observations from that season's playoffs.

[11] We drop observations for the last shot of a set (usually the second shot of two) with the shooter's team losing by two or three points with five or fewer seconds remaining, as players sometimes intentionally miss their shot in those situations in the hope that their team will rebound the missed shot. We examined the data closely, and found no evidence of these intentional misses for shots taken with greater than five seconds remaining.

1 is that better overall shooters are more likely to be selected to shoot when the shooter's team is winning by at least one point, as compared to when his team is tied or losing, especially in the last 15 seconds. That is, the "Normal" plot for the bottom right figure, which represents the overall free throw percentage of players shooting for each score difference, is approximately flat for score differences between -5 and 0, and also for score differences between 1 and 5. But the plot jumps when the difference increases from 0 to 1. This also occurs to some extent when there are 16-30 seconds remaining, and likely reflects the fact that teams that are winning in the final seconds know they are likely to be fouled, and as a result attempt to give the ball to their players known to be relatively good shooters, and perhaps even players thought to perform especially well under pressure. This may bias our estimates of choking towards zero for situations in which the shooter's team is winning in the final 30 seconds; we will keep this in mind as we proceed.

Our empirical strategy is to use linear probability models to estimate the effects of several variables that may be associated with pressure on a binary dependent variable that equals one if the free throw is made (*FreeThrowMade*). We use player-season fixed effects, which means our estimates should be interpreted as deviations from player-season means, and should eliminate many possible selection biases (shooters who are better—or worse—on average may be more likely to get fouled in certain situations).[12] However, our results may still be affected by selection (changes in the composition of shooters in different situations), as we discuss. We also include quarter dummies and dummies for whether the shot was from a one-shot set (*OneShot*, which usually occurs when the player was fouled during a successful field goal attempt), and whether the previous free throw was made (*PrevMade*) and whether it was missed (*PrevMissed*) for the second or third free throw of a set, as controls (see Arkes, 2010).

We use linear models rather than logit or probit because most of our estimates of interest are coefficients on interaction terms, which are difficult to interpret for non-linear models (Ai and Norton, 2003). It is well known that the linear probability model generally yields results that are very similar to those of logit and probit, especially for dependent variables whose means are not too close to zero or one (see, e.g., Angrist and Pischke, 2009), and we have confirmed that

[12] We experimented with alternative controls, including player fixed effects, and player 60-120 day moving averages, to allow for players' free throw percentage to vary less/more frequently than season to season, and obtained similar results.

results are very similar when we use logit and calculate marginal effects using the methods of Ai and Norton (2003) and Corneliben and Sonderhof (2009).[13]

We first estimate the effects of several variables (home/playoff status; attendance, score differences) on performance in the first three quarters of games, to examine the effects of factors that may affect anxiety throughout games. We drop the fourth quarter to ensure that our results are not affected by end-of-game effects. The theoretical effect of home status on performance here is ambiguous, as players may benefit from the support of fans at home, but may also feel more pressure to please them. Clearly, playoff games and games with higher attendance levels should involve higher pressure, even in the first three quarters. Smaller score differences should involve higher pressure, but the differences may not be substantial, since in the first three quarters of a game it is highly unlikely that any single free throw will affect the game's outcome (and it is almost certain that fans and analysts will not attribute the game's outcome to a free throw taken before the fourth quarter).

We then analyze how the effects of the score difference variables may change in the final seconds when the game's outcome is on the line. To do this we interact dummies for various score differences with dummy variables for whether there is less than 60/30/15 seconds (*Last60, Last30, Last15*) left in the game (fourth quarter or overtime). We use interactions so that we can include the *Last* variables separately to control for fatigue and other changes that may result at the end of games for reasons other than pressure. We use dummy variables for different amounts of time remaining to transparently allow for a non-linear relation between score difference and pressure. We focus on the last minute since, as shown in Figure 1, there is little evidence of choking with greater than one minute remaining.

We also estimate coefficients on a number of triple-interaction terms involving score difference, the *Last* variables, and other variables. First, we look at interactions with *FinalShot*, a dummy for whether the free throw is the last of its set (usually second of two). In many situations we would expect pressure to be higher when the shooter is down to his last shot; for example, if the shooter's team is up one point, then there is a large benefit to hitting one shot, since that puts his team up by a field goal, and the benefit of hitting the second shot is relatively low. Thus, pressure is likely much higher when *FinalShot*=1, for that score difference. We also

---

[13] Results are not reported, but available on request. Another reason that we prefer linear models is that due to our large sample size and player-season fixed effect specification it is computationally extremely demanding to use the Ai and Norton (2003) approach.

examine interactions with shooter characteristics: free throw percentage for the season for shots taken with more than one minute remaining in the fourth quarter or overtime (*FTPct*[14]), years of experience (*Exper*), number of attempts per game for the season (*Atts*) and number of attempts for the season in pressure situations (*PrAtts*), defined as attempts in the final minute when the shooter's team is tied or losing by up to four points. These interactions allow us to analyze the types of players who are more or less likely to choke.[15] The estimate for *PrAtts* may be most difficult to interpret. The variable is intended to proxy experience shooting free throws in high pressure situations. If its coefficient is positive, this could be caused by selection (more clutch players may be given the ball more often in pressure situations) or could indicate that greater experience in the clutch causes clutch performance to improve. Thus, if we find the estimate is positive, it is not clear which of these factors would be responsible. On the other hand, if we find the estimate is insignificant, we could view this as evidence that selection does not substantially affect our overall results, at least for situations in which the shooter's team is losing or tied. We exclude situations in which the shooter's team is winning by a small number of points from the definition of this variable because selection effects are likely more severe in those situations, since the shot is much more likely to be due to a strategically committed foul.[16]

The coefficients on the score dummies can be interpreted as the difference between the expected value of the dependent variable when the score dummy equals one, and the expectation when the score difference is in the omitted category. Thus, the definition of the omitted category affects the interpretation of the results. To make results as transparent as possible we present results with two "baselines" (omitted categories): 1) down or up by 11+ points (11+ baseline), 2) down or up by five to 10 points (5-10 baseline). The 11+ baseline is intended to capture the

[14] We exclude the current observation in the calculation of this variable so that it is not endogenous. This exclusion is clearly only relevant for observations that occur with greater than one minute remaining in the game.
[15] In all specifications with interaction terms we include all variables used in the interaction terms as separate regressors. We drop the player-season fixed effects in specifications involving *FTPct* to avoid perfect multicollinearity. The variable *FTPct* still provides a very good control for overall free throw shooting ability.
[16] As discussed above, Figure 1 shows that better shooters are more likely to get the ball when teams are winning by a small margin, but not when losing. We do not think possible sample selection biases are too serious a concern for situations in which teams are losing by a small margin since basketball teams tend to simply give the ball to their best shooters in these situations, and not the players thought to be "most clutch" since this is so uncertain. However, we acknowledge that selection may affect our results to some extent. They would likely bias our estimates of performance under pressure upward (if viewed as estimates of average clutch performance across all players), since the shooting team has more control over the foul shooter than the defending team. Since the direction of this bias goes against our main conclusion—that players do choke—we think that, qualitatively at least, our main conclusion is on safe ground. We can also view our estimates as lower bounds on average choking, given the direction of possible selection bias.

lowest pressure situations, and the 5-10 baseline represents medium pressure situations. We combine situations in which a team is either winning or losing by the same range of points for each baseline because pressure, as determined by the expected effect of the shot on game outcome, should be similar in both cases. This approach allows us to more easily capture a non-linear (for example, inverse-U) relation between pressure and performance. The approach also allows us to test whether performance changes significantly when the score is very close as compared to both baselines (if we just used the 11+ baseline we would have to use F-tests to test the difference between estimates for small score margins, and margins of five to 10 points), and to check that our results are not driven by changes in the composition of players that occur when the score difference takes different values.[17]

4. Results

Table 2 presents results for the analysis of the first three quarters. The magnitude of the coefficients on the score difference are less than or equal to 1% and mostly insignificant, so we do not report them. There is a small positive effect to being the home team, and small negative effect of higher attendance, both of which are only significant when the data from all three quarters are pooled. The attendance effect is not significantly different for home or away teams, and playoff status does not have an economically or statistically significant effect. The estimates are generally similar for each of the first three quarters. Almost all of the effects are not significantly affected by player characteristics, but there is some evidence that the playoff effect is increasing in overall shooting percentage, but decreasing in player experience (the coefficient on *Playoff*FTPct* is positive and significant at 10%, and on *Playoff*Exper* is negative and significant). The former is consistent with more confident players choking less in higher pressure situations, but the latter result is difficult to explain.

The paper's main results are presented in Tables 3 and 4. Both present estimated coefficients for interactions of dummies for various score differences with the *Last* variables, and

[17] It is well known that teams are more likely to play "scrubs" when losing or winning by a substantial margin at the end of the game. Scrubs may actually feel even less pressure than normal players in these situations (since they likely know they are not being evaluated by their coach and others as closely as they are evaluated in the middle of games), and will not be as fatigued, which could bias the estimates on the *Last* variables upward, and consequently the estimate on *Last* interacted with other score variables to be biased downward. We are confident our results are robust to this issue since results are similar when we use the 5-10 baseline, and we also found results are similar for the 11+ baseline when we drop scrubs (players who took most of their free throw attempts in the last minute when winning or losing by 11+ points).

in separate specifications, dummies for whether the shot was the final shot from the set. In Table 3 the omitted score category is winning or losing by 11 or more points, and in Table 4 the omitted score difference is 5-10 points. The non-interacted score variables and *Last* variables' estimates are unreported as they are always less than 1% and usually insignificant except for the model with *Last15* in Table 4 (the coefficient on *Last15* is -1.8% with p-value 0.011, indicating players shoot worse in the last 15 seconds, given the score difference is 5-10). Thus, the estimates in Table 3 can be interpreted as the approximate difference in percentage when the interaction term equals one, as compared to when the score difference is greater than 10 at any time in the game. The table shows players shoot 4.8% and 6.6% worse when down one or two, in the last 30 seconds; the respective numbers for the last 15 seconds are 6.8% and 8.8%. Players shoot around 5% worse when up one or two in the last 15 seconds on the final shot of a set. This is consistent with the choking hypothesis, since pressure should be higher on the final shot. On the other hand, players perform better on the initial shot of a set, as compared to the final shot, when down by one point, which is puzzling. The estimates for interactions of dummies for down three and down four with *Last30* and *Last15* range from -1% to -5%, but are statistically insignificant and unreported. Table 4 shows that players shoot 3.7% and 5.5% worse when down one or two in the last 30 seconds, as compared to when the score difference is five to 10. The numbers are 4.5% and 6.5%, respectively, in the last 15 seconds, though only the latter is significant. Players again shoot around 5% worse on the final shot in the last 15 seconds when winning by one or two. These effects are smaller than those reported in Table 3 because players shoot 1-2% worse in the last 30 seconds, as compared to the rest of the game, when the score difference is five to 10, indicating there are small pressure effects for these moderate score differences. The estimates tend to be of much lower magnitude and insignificant for the *Last60* models, and for interactions involving the dummy for the score being tied.[18]

Table 5 presents estimates of interaction terms involving player characteristics, for the +11 baseline.[19] We create a new score difference dummy variable for this purpose, *Down1_4*, equal to one when the shooter's team is down by one to four points, as Tables 3 and 4 all indicate

[18] Figure 1 indicates that players perform worse when the score is tied when there are 16-30 seconds remaining; however, we find even these effects are insignificant when we test for them using this framework.

[19] In the interest of brevity we do not present results for the alternate baselines for the remaining tables, as they are largely similar. We also do not report results for *Up0_4* interactions, although many are significant. It is worth noting in particular that the interaction with *FTPct* is also significant, but of a lower magnitude than the *Down1_4* estimate, which is consistent with choking being more likely when teams are losing by a small margin rather than winning for players of all abilities.

that choking effects are greater when the shooter's team is losing by a small margin, rather than tied or winning and Figure 1 and unreported regression results indicate performance does decline when losing by small margins greater than two (results are similar when we use different versions of this variable, such as down or up by one to four or one to two). For this table many of the coefficients unreported to save space are significant and large, so we need to be careful with interpretation. The estimate for *FTPct* interacted with *Down1_4* and the *Last* variables is approximately 0.5-0.6 and highly significant, implying that for every 10% improvement in *FTPct*, the effect of *Last* on the effect of *Down1_4* declines by 5-6%. Since this cross-partial effect is our primary measure of choking, and *FTPct* does vary substantially across players (approximately from 50% to 90%), this implies substantial heterogeneity in the degree of choking. Total number of attempts in the season and years of experience have insignificant effects. Number of attempts in pressure situations in that season (*PrAtts*) is insignificant as well, which should alleviate sample selection concerns (since if more clutch performers got more attempts in the clutch, the estimates for *PrAtts* should be positive). Results are similar whether the interactions are included in separate regressions or all together.

Table 6 presents results for additional interactions with *Down1_4*. Models 3 and 4 exclude the attendance variable as it is not available for the last two seasons of our sample, and for numerous observations for the earlier seasons. The most striking result is that performance in the final 30 seconds when *Down1_4* equals one improves by around 5% when the previous free throw was made, and declines by over 5% when the previous shot was missed, though these effects are not estimated precisely. Home status and attendance have no significant interaction effects (the direction of the home effect is positive but its magnitude is always less than 4%). Playoff status does not interact significantly with *Down1_4* and *Last*, and the magnitudes of the effects are close to zero for models 3 and 4. We thought this may be due to players gaining experience throughout the season, and consequently becoming less prone to choke later in the season, which could nullify increased choking in the playoffs. This does not appear to be the case, however, as when we include an interaction with a variable equal to month of season, which is highest in playoff months and should thus pick up the effect of declining within-season choking, results are similar (results not reported). We recognize that pressure may also be relatively high for games that occur late in the regular season that are likely to affect playoff seeds. To account for this we estimate the models on a sub-sample without regular season games played in the final month (April); this should accentuate the difference in pressure between playoff and regular season games. Doing this does not substantially change the *Playoff* results, but does weaken the *PrevMade* and *PrevMiss* estimates, so we do not report these results.

5. Discussion and Conclusion

We have found strong evidence that NBA players do choke, and to a substantial degree, but only at the very end of games when their team is winning or losing by a small margin. Choking effects are stronger for players whose teams are losing, which is unsurprising since the marginal effect of a foul shot on the probability of winning the game is likely greater for shooters whose teams are losing rather than winning. However, the choking estimates for players on winning teams may be biased slightly downward, as the shooters in those situations are slightly better than average (see Table 1) and we have evidence that better shooters are less likely to choke. We do find it surprising that players do not choke when the score is tied in the last 15 seconds, given the other results. Perhaps this is because players perceive that even if the shot is missed the shooter cannot be held directly responsible for the game outcome—the shooter's opponent still needs to score in the final seconds to win, and if the opponent scores by making a two or three point field goal (as scoring usually occurs) a single free throw would not make a difference to the game outcome anyway. On the other hand, when the shooter's team is up by one or two, a single free throw could be the difference between a two or three point shot by the opponent leading to a win or overtime. However, these perceptions would be at odds with the empirical reality. We find that when games are tied in the last 15 seconds, after a made free throw the shooter's team has a 69.5% chance of winning in the current period, and only a 35.0% chance of winning after a miss (after a miss teams are approximately equally likely to lose or go to an extra period). When up one point, the percentages are 71.1% and 67.7%, respectively. The numbers are similar for teams up two points. This indicates that the shot is very important to the team's outcome when the game is tied, but is unlikely to matter when the shooter's team is winning by a small margin.[20]

We are also somewhat surprised that home and playoff status do not have substantial effects on choking (despite the ambiguous theory regarding home status, discussed above).

[20] The percentages are 28.7% and 5.5% when down one point, respectively. The standard errors are all less than 5%. These data are only available for fourth quarters of games from the 02-03 through 07-08 seasons.

Perhaps the lack of a home effect is due to players knowing their teammates, coach, and the vast majority of fans, will be aware of or learn about the performance (good or bad) whether it takes place at home or on the road, and so the effects of choking on the player's reputation are independent of home status. The lack of a home effect may also be due to competing forces; players may feel more pressure playing at home, but be more distracted by opposing fans on the road. The lack of a playoff effect may be due to small sample size; although our dataset is large, the number of observations for very high pressure situations that occurred in the playoffs is still highly limited.

NBA players are highly-paid professionals, most of whom have substantial experience performing in front of thousands watching in person, and millions watching on television. Our findings that even NBA players choke lead us to believe that athletes across all sports, and workers in general, are likely subject to anxiety in pressure situations that cause them to perform less than optimally. Topics for future research include deeper analysis of player-level heterogeneity and consistency of player-level clutch performance across seasons, and the relation between performance under pressure and earnings.

References

Angrist, J.D. and J.S. Pischke. 2009. *Mostly Harmless Econometrics*, Princeton University Press.

Ariely, Dan. "The Irrational Side of Corporate Bonuses." 27 May 2010 <www.huffingtonpost.com>.

Ariely, D., Gneezy, U., Loewenstein, G. and Mazar, N. 2009. "Large Stakes and Big Mistakes." *Review of Economic Studies*, 76, 2, 451—469

Arkes, Jeremy. 2010. "Revisiting the Hot Hand Theory with Free Throw Data in a Multivariate Framework." *Journal of the Quantitative Analysis in Sports*, 6, 1, 2.

Bannier, C.E. and E. Feess. 2010. "When High-Powered Incentive Contracts Reduce Performance: Choking Under Pressure as a Screening Device." *Frankfurt School- Working Paper Series*.

Corneliben, T. and K. Sonderhof. 2009. "Partial Effects in Probit and Logit Models with a Triple Dummy Variable Interaction Term." *The Stata Journal*, 9, 4, 571-583.

Dohmen, T.J. 2008. "Do Professionals Choke Under Pressure?" *Journal of Economic Behavior and Organization*, 65, 3-4, 636-653.

Hilla, D.M., S. Hantonb, N. Matthewsa, and S. Fleming. 2010. “Choking in Sport: A Review.” *International Review of Sports and Exercise Psychology*, 3, 1, 24-39.

Levitt, S.D. and J.A. List. 2007. “What do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?” *The Journal of Economic Perspectives*, 153-174.

Neyer, Rob. “Clutch Hitting.” 17 Dec. 1999 <www.espn.com>.

Ruah, M.T. and G. Seccia. 2010. “Agency and Anxiety.” *Journal of Economics and Management Strategy*, 19, 1, 87-116.

Wallace, H.M., R.F. Baumeister, and K.D. Vohs. 2005. “Audience Support and Choking Under Pressure: A Home Disadvantage?” *Journal of Sports Sciences*, 23, 4, 429-438.

Wang, Jin and Marchant, Daryl and Morris, Tony (2004). “Coping styles and susceptibility to choking.” *Journal of sport behavior*, 27 (1). pp. 75-92.

Worthy, D.A., A.B. Markman, and W.T. Maddox. 2009. “*The International Journal of Creativity and Problem Solving*, 19, 1, 53-58.

Yarett, Ian. “The Psychology of Missed Field Goals: Was Nate Kaeding’s Performance Part of a Choking Outbreak?” 22 Jan. 2010 <www.newsweek.com>.

Figure 1

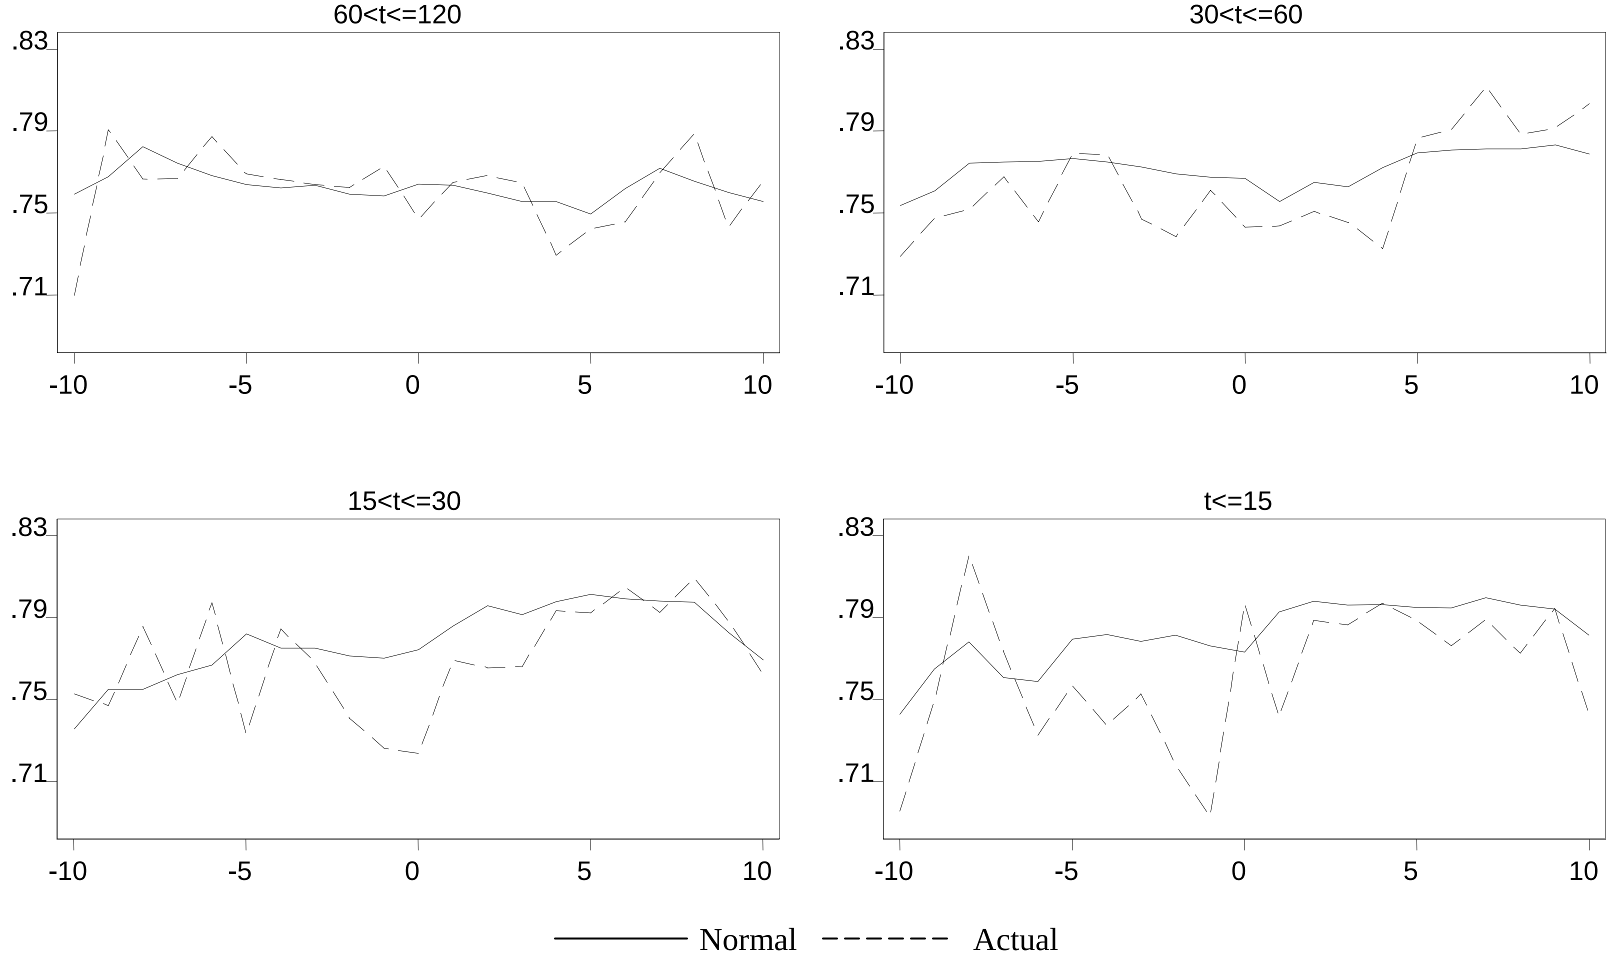

Notes*:* *t* = seconds remaining in fourth quarter or overtime; Normal = mean *FTPct*; Actual = mean *FreeThrowMade*; X-axis = shooter’s team’s score minus opponent’s score, prior to shot.

Table 1: Summary Statistics

	t>60			30<t<=60			15<t<=30			t<=15		
	mean	std error	N	mean	std error	N	mean	std error	N	mean	std error	N
Panel A: Summary Statistics for *FreeThrowMade*												
Home = 1												
Down1_4	0.750	0.002	35323	0.769	0.017	585	0.739	0.024	341	0.729	0.021	468
Up0_4	0.750	0.002	48831	0.755	0.015	867	0.774	0.011	1524	0.778	0.008	3007
Other	0.752	0.001	135517	0.774	0.006	4423	0.779	0.007	3298	0.780	0.008	2431
Home = 0												
Down1_4	0.751	0.002	36632	0.738	0.017	682	0.756	0.022	386	0.722	0.021	454
Up0_4	0.751	0.002	42503	0.728	0.016	767	0.776	0.012	1128	0.791	0.008	2498
Other	0.750	0.001	131434	0.760	0.006	4546	0.779	0.008	2876	0.771	0.009	2079
Playoff = 1												
Down1_4	0.742	0.007	4447	0.816	0.042	87	0.830	0.052	53	0.687	0.057	67
Up0_4	0.744	0.006	5741	0.738	0.047	88	0.735	0.035	162	0.755	0.022	376
Other	0.753	0.003	17357	0.765	0.016	686	0.770	0.023	330	0.728	0.026	301
Panel B: Summary Statistics for Other Variables												
FTPct												
Down1_4	0.752	0.000	76402	0.771	0.002	1354	0.772	0.003	780	0.779	0.003	989
Up0_4	0.751	0.000	97075	0.765	0.002	1722	0.793	0.001	2814	0.795	0.001	5881
Other	0.751	0.000	284308	0.762	0.001	9655	0.778	0.001	6504	0.782	0.001	4811
Exper												
Down1_4	5.201	0.013	76402	5.375	0.094	1354	5.031	0.123	780	5.332	0.114	989
Up0_4	5.234	0.012	97075	5.382	0.085	1722	5.550	0.066	2814	5.473	0.047	5881
Other	4.997	0.007	284308	4.758	0.037	9655	4.908	0.045	6504	4.928	0.052	4811
Atts												
Down1_4	3.763	0.009	76402	4.281	0.070	1354	4.139	0.087	780	4.182	0.086	989
Up0_4	3.808	0.008	97075	3.926	0.061	1722	3.702	0.045	2814	3.648	0.031	5881
Other	3.586	0.005	284308	3.193	0.025	9655	3.098	0.028	6504	3.159	0.033	4811
PrAtts												
Down1_4	2.822	0.013	76402	5.756	0.108	1354	5.609	0.136	780	5.657	0.125	989
Up0_4	2.816	0.011	97075	3.686	0.097	1722	3.460	0.074	2814	3.243	0.049	5881
Other	2.625	0.006	284308	2.517	0.035	9655	2.514	0.041	6504	2.592	0.053	4811

Notes: The first two sections of panel A (*Home* = 1 and *Home* = 0) give the statistics for regular season games only. Observations for *OneShot*, second and third shot with team losing by 2 or 3 points and time remaining is less than 6 second are dropped. t = seconds remaining in fourth quarter or overtime.

Table 2: Analysis of Quarters 1-3

	(1) quarter 1	(2) quarter 2	(3) quarter 3	(4) quarter 1-3	(5) quarter 1	(6) quarter 2	(7) quarter 3	(8) quarter 1-3
Home	0.002	0.004	0.005	0.004**	-0.034	0.026	0.003	0.001
	[0.003]	[0.003]	[0.003]	[0.002]	[0.027]	[0.024]	[0.024]	[0.014]
*Home*FTPct*					0.052	-0.039	0.016	0.007
					[0.035]	[0.032]	[0.031]	[0.019]
*Home*Exper*					0.000	-0.000	-0.001*	-0.001
					[0.001]	[0.001]	[0.001]	[0.000]
*Home*Atts*					-0.001	0.003**	-0.001	0.000
					[0.001]	[0.001]	[0.001]	[0.001]
Attend	-0.013	-0.012	-0.008	-0.012***	-0.008	-0.016	-0.099	-0.044
	[0.009]	[0.008]	[0.008]	[0.005]	[0.072]	[0.063]	[0.060]	[0.038]
*Attend*FTPct*					0.023	-0.024	0.133*	0.045
					[0.096]	[0.084]	[0.080]	[0.050]
*Attend*Exper*					-0.003	0.004*	-0.002	-0.000
					[0.002]	[0.002]	[0.002]	[0.001]
*Attend*Atts*					-0.002	-0.000	-0.000	-0.001
					[0.004]	[0.003]	[0.003]	[0.002]
*Home*Attend*	-0.005	0.018	-0.016	-0.001	0.058	-0.025	-0.025	-0.002
	[0.013]	[0.012]	[0.012]	[0.007]	[0.097]	[0.086]	[0.084]	[0.052]
*Home*Attend*FTPct*					-0.085	0.084	0.012	0.015
					[0.129]	[0.115]	[0.111]	[0.069]
*Home*Atten*Exper*					0.002	-0.003	0.001	-0.000
					[0.003]	[0.003]	[0.003]	[0.002]
*Home*Atten*Atts*					-0.001	-0.002	0.001	-0.001
					[0.005]	[0.005]	[0.005]	[0.003]
Playoff	0.006	-0.006	0.008	0.004	-0.019	-0.047	0.011	-0.020
	[0.007]	[0.006]	[0.006]	[0.004]	[0.052]	[0.048]	[0.045]	[0.028]
*Playoff*FTPct*					0.096	0.087	0.000	0.060*
					[0.065]	[0.060]	[0.058]	[0.036]
*Playoff*Exper*					-0.006***	-0.004**	-0.000	-0.003***
					[0.002]	[0.002]	[0.002]	[0.001]
*Playoff*Atts*					-0.003	-0.000	-0.001	-0.001
					[0.002]	[0.002]	[0.002]	[0.001]
N	67,585	87,472	89,413	244,470	67,585	87,472	89,413	244,470
R-squared	0.003	0.003	0.002	0.003	0.043	0.042	0.045	0.043

Notes: All models include *OneShot, PrevMade, PrevMiss, Up5_10, Up4, Up3, Up2, Up1, Tied, Down1, Down2, Down3, Down4, Down5_10.* Model 4&8 also include dummies for quarter. *Attend* is measured in ten thousands. Models 1-4 include player-season fixed effects, and models 5-8 include *FTPct*, *Exper* and *Atts* as separate regressors. Robust standard errors given in brackets. *, **, *** denote 10%, 5% and 1% significance.

Table 3: Analysis of Last Minute Pressure Effects, 11+ Baseline

	(1) X=60	(2) X=30	(3) X=15	(4) X=60	(5) X=30	(6) X=15
*LastX*Down2*	-0.037**	-0.045**	-0.063**	-0.021	-0.034	-0.060
	[0.017]	[0.023]	[0.032]	[0.024]	[0.031]	[0.042]
*LastX*Down1*	-0.041**	-0.066***	-0.088***	-0.062**	-0.097***	-0.108**
	[0.016]	[0.022]	[0.030]	[0.024]	[0.032]	[0.043]
*LastX*Tied*	-0.015	-0.008	0.018	-0.011	-0.010	0.023
	[0.017]	[0.022]	[0.029]	[0.025]	[0.032]	[0.040]
*LastX*Up1*	-0.026**	-0.029**	-0.042**	-0.009	-0.015	-0.016
	[0.012]	[0.015]	[0.020]	[0.015]	[0.017]	[0.023]
*LastX*Up2*	-0.018*	-0.019	-0.018	0.001	0.002	0.008
	[0.011]	[0.013]	[0.019]	[0.014]	[0.017]	[0.022]
*FinalShot*LastX*Down2*				-0.029	-0.020	0.001
				[0.032]	[0.042]	[0.057]
*FinalShot * LastX*Down1*				0.043	0.060	0.042
				[0.031]	[0.041]	[0.053]
*FinalShot * LastX*Tied*				-0.006	0.005	-0.009
				[0.032]	[0.041]	[0.050]
*FinalShot * LastX*Up1*				-0.036	-0.031	-0.063**
				[0.023]	[0.026]	[0.032]
*FinalShot * LastX*Up2*				-0.034*	-0.039**	-0.046**
				[0.019]	[0.020]	[0.023]
N	436,898	424,167	414,069	436,898	424,167	414,069
R-squared	0.003	0.003	0.003	0.003	0.003	0.003

Notes: All models include *Up5_10, Up4, Up3, Up2, Up1, Tied, Down1, Down2, Down3, Down4, Down5_10, Last60/30/15, PrevMade, PrevMiss Home, Playoff*, interaction of *Last60/30/15* and score dummies, dummies for quarter. Models 1-3 also include *OneShot*. Models 4-6 also include *FinalShot* and its separate interactions with *Last60/30/15* and score dummies. Observations with *FinalShot*=1 and less than 6 seconds remaining (in fourth quarter or overtime) are dropped, as are observations with *LastX*=0 with less than five minutes remaining (this is why N varies from model to model). All models include player-season fixed effect and use robust standard errors are given in brackets. *, **, *** denote 10%, 5% and 1% significance.

Table 4: Analysis of Last Minute Pressure Effects, 5-10 Baseline

	(1) X=60	(2) X=30	(3) X=15	(4) X=60	(5) X=30	(6) X=15
*LastX*Down2*	-0.033**	-0.035	-0.041	-0.018	-0.025	-0.039
	[0.016]	[0.021]	[0.029]	[0.024]	[0.030]	[0.039]
*LastX*Down1*	-0.037**	-0.055***	-0.065**	-0.059**	-0.087***	-0.087**
	[0.016]	[0.020]	[0.027]	[0.024]	[0.031]	[0.040]
*LastX*Tied*	-0.010	0.003	0.040	-0.008	-0.000	0.045
	[0.016]	[0.021]	[0.026]	[0.024]	[0.031]	[0.037]
*LastX*Up1*	-0.022*	-0.019	-0.019	-0.005	-0.005	0.005
	[0.011]	[0.012]	[0.015]	[0.014]	[0.015]	[0.018]
*LastX*Up2*	-0.013	-0.008	0.005	0.004	0.012	0.029
	[0.010]	[0.010]	[0.013]	[0.014]	[0.015]	[0.018]
*FinalShot*LastX*Down2*				-0.029	-0.019	0.002
				[0.032]	[0.042]	[0.057]
*FinalShot * LastX*Down1*				0.042	0.060	0.043
				[0.031]	[0.041]	[0.053]
*FinalShot * LastX*Tied*				-0.007	0.005	-0.008
				[0.032]	[0.041]	[0.050]
*FinalShot * LastX*Up1*				-0.036	-0.031	-0.062**
				[0.023]	[0.026]	[0.032]
*FinalShot * LastX*Up2*				-0.034*	-0.039*	-0.046*
				[0.019]	[0.020]	[0.023]
N	436,898	424,167	414,069	436,898	424,167	414,069
R-squared	0.003	0.003	0.003	0.003	0.003	0.003

Notes: All models include *Up5_10*, *Up4*, *Up3*, *Up2*, *Up1*, *Tied*, *Down1*, *Down2*, *Down3*, *Down4*, *Down5_10*, *Last60/30/15*, *PrevMade*, *PrevMiss Home*, *Playoff*, interaction of *Last60/30/15* and score dummies, dummies for quarter. Models 1-3 also include *OneShot*. Models 4-6 also include *FinalShot* and its separate interactions with *Last60/30/15* and score dummies. Observations with *FinalShot*=1 and less than 6 seconds remaining (in fourth quarter or overtime) are dropped, as are observations with *LastX*=0 with less than five minutes remaining (this is why N varies from model to model). All models include player-season fixed effect and use robust standard errors are given in brackets. *, **, *** denote 10%, 5% and 1% significance.

Table 5: Player Characteristics Interactions

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: *Last30* Interactions						
*Last30*Down1_4*	-0.038***	-0.453***	-0.075***	-0.057**	-0.071***	-0.475***
	[0.014]	[0.121]	[0.025]	[0.023]	[0.021]	[0.122]
*FTPct*Last30*Down1_4*		0.557***				0.547***
		[0.155]				[0.159]
*Atts*Last30*Down1_4*			0.008			-0.002
			[0.006]			[0.008]
*Exper*Last30*Down1_4*				0.004		0.001
				[0.004]		[0.004]
*PrAtts*Last30*Down1_4*					0.007	0.003
					[0.004]	[0.006]
N	424,167	424,167	424,167	424,167	424,167	424,167
R-squared	0.003	0.04	0.003	0.003	0.003	0.04
Panel B: *Last15* Interactions						
*Last15*Down1_4*	-0.058***	-0.511***	-0.093***	-0.045	-0.086***	-0.546***
	[0.021]	[0.189]	[0.035]	[0.034]	[0.031]	[0.193]
*FTPct*Last15*Down1_4*		0.611***				0.648***
		[0.243]				[0.250]
*Atts*Last15*Down1_4*			0.008			0.008
			[0.009]			[0.013]
*Exper*Last15*Down1_4*				-0.004		-0.007
				[0.006]		[0.006]
*PrAtts*Last15*Down1_4*					0.002	-0.008
					[0.007]	[0.010]
N	414,069	414,069	414,069	414,069	414,069	414,069
R-squared	0.003	0.04	0.003	0.003	0.003	0.04

Notes: All models include *Up5_10, Up0_4, Close, Down5_10, Last30/15, Oneshot, PrevMade, PrevMiss, Home, Playoff,* interaction of *Last30/15* and score dummies, dummies for quarter. Models with *FTPct*, *Atts*, *Exper* and *PrAtts* interactions include these variables separately and separate interactions with *Last30/15* and score dummies. Observations with *FinalShot*=1 and less than 6 seconds remaining (in fourth quarter or overtime) are dropped, as are observations with *LastX*=0 with less than five minutes remaining (this is why N varies from model to model). All models except those 2 and 6 include player-season fixed effects. Robust standard errors are given in brackets. *, **, *** denote 10%, 5% and 1% significance.

Table 6: Other Interactions

	(1) X=30	(2) X=15	(3) X=30	(4) X=15
*PrevMade*LastX*Down1_4*	0.042	0.058	0.049**	0.045
	[0.026]	[0.036]	[0.023]	[0.031]
*PrevMiss*LastX*Down1_4*	-0.036	-0.046	-0.069*	-0.096**
	[0.041]	[0.053]	[0.035]	[0.047]
Home LastX*Down1_4*	0.016	0.032	0.004	0.022
	[0.025]	[0.034]	[0.022]	[0.029]
Playoff LastX*Down1_4*	-0.046	-0.108	0.036	-0.003
	[0.053]	[0.072]	[0.043]	[0.060]
Attend LastX*Down1_4*	0.012*	0.005		
	[0.007]	[0.009]		
*Home*Attend*LastX*Down1_4*	-0.008	-0.004		
	[0.009]	[0.013]		
N	308,176	300,664	424,167	414,069
R-squared	0.003	0.003	0.003	0.003

Notes: All models include *Up5_10, Up0_4, Close, Down5_10, Last30/15, Oneshot, PrevMade, PrevMiss, Home, Playoff,* interaction of *Last30/15* and score dummies, dummies for quarter, separate interaction of *Home* and *Playoff* with *Last30/15* and score dummies. Models 1-2 also include *Attend* and *Home*Attend* and their separate interactions with *Last30/15* and score dummies. Only observation from 02-03 to 07-08 season are used for models 1-2. Observations with *FinalShot*=1 and less than 6 seconds remaining (in fourth quarter or overtime) are dropped, as are observations with *LastX*=0 with less than five minutes remaining (this is why N varies from model to model). All models include player-season fixed effect and use robust standard errors are given in brackets. *, **, *** denote 10%, 5% and 1% significance.