HEAT CHECK: NEW EVIDENCE ON THE HOT HAND IN BASKETBALL*

Andrew Bocskocsky

John Ezekowitz

Carolyn Stein†

July 8, 2014

Abstract

The vast literature on the hot hand fallacy in basketball rests on the assumption that shot selection is independent of player-perceived hot or coldness. In this paper, we challenge the assumption of independence using a novel dataset of over 83,000 shots from the 2012-2013 NBA season, combined with optical tracking data of both the players and the ball. We use this data to show that players who have exceeded their expected shooting percentage over recent shots shoot from significantly further away, face tighter defense, are more likely to take their team's next shot, and overall attempt more difficult shots. In other words, we show that the independence assumption fails. We then turn to the hot hand itself and show that players who are outperforming (i.e. are "hot") are more likely to make their next shot if we control for the difficulty of that shot. We estimate a 1.2% increase in the likelihood of the typical player making his next shot for each additional prior shot he made. JEL codes: D83, D84, L83.

*We would like to thank Roland Fryer, Matt Davis, Larry Katz, and Carl Morris for their help and insight. We also thank Brian Kopp, Ryan Warkins, and STATS, Inc. for providing us with the raw data source. All remaining errors are our own.

†*Address:* 658 Massachusetts Avenue Apartment 1B, Boston MA 02118. *Fax:* (617)-516-2710 *Telephone:* (781)-572-7726 *Email:* carolynstein@post.harvard.edu

1 Introduction

Humans are inherently bad at handling uncertainty and chance. We crave order and symmetry, and this tendency often leads us astray in the face of randomness. This is evidenced by the belief in "local representativeness," or the "law of small numbers," first formally defined by Kahneman and Tversky (1971). Humans, the authors write, "expect the essential characteristics of a chance process to be represented not only globally in the entire sequence, but also locally, in each of its parts" (1985). For example, we tend to overestimate the likelihood that a sequence of four coin flips contains two heads, because we want this short sequence to "represent" the overall 50/50 nature of the coin flipping process. A series of four consecutive heads might cause us to question if a coin is truly fair, because we underestimate the probability of such an event happening by chance. As Professor Larry Summers put it in a recent discussion with the Harvard Men's Basketball team, "people apply patterns to randomness" (Davidson, 2013). The four consecutive coin flips are random, but we are inclined to attribute a cause or "pattern" - in this case a biased coin - to explain the statistical noise.

The "hot hand" or "streak shooting" has also been used as an example of local representativeness. An average NBA viewer observes Player X, a sixty percent shooter, make four shots in a row. He thinks to himself, "this isn't representative of his normal sixty percent shooting – Player X must be hot!" The statistician, however, knows that streaks of four makes in a row are inevitable over the course of a season. Short sequences of shots need not be consistent with the player's overall shooting percentage.

So who is correct – the NBA fan or the statistician? It depends on the frequency of these streaks. If there are too many – more than would be expected in series of independent shots – then the viewer is correct. However, if the number of observed streaks is consistent with the number expected from a series of independent, identically distributed Bernoulli trials, then the statistician is correct. This debate, seemingly long-ago settled in favor of the statistician, has given rise to perhaps the most well-known theories in behavioral economics: the hot hand fallacy.

The seminal piece of research that attempted to answer this question was done by Thomas Gilovich, Robert Vallone, and Amos Tversky (1985). In this paper, the authors analyzed series of shots by players on the Philadelphia 76ers, looking for a positive correlation between successive shots, and found none. They also analyzed series of free throws by the Boston Celtics and the Cornell men's and women's basketball teams, and again found no evidence of serial correlation. Subsequent studies, including Adams, 1992; Koehler and Conley, 2003; Bar-Eli, Avugos and Raab, 2006; and Rao 2009 have confirmed this finding. In the recent academic literature, only Arkes (2010) has found a significant hot hand effect in free throw shooting in the 2005-2006 season.

Today, among the academic crowd, the hot hand in live game action is almost universally considered a "fallacy." In a recent op-ed, David Brooks cites the hot hand and the original Gilovich, Vallone, and Tversky paper as a "really good [example of] exposing when our intuitive vision of reality is wrong" (2013). Yet, among basketball fans and players, the hot hand is a myth that refuses to die. Kahneman and Tversky found that among college basketball fans, 91% believed that a player has "a better chance of making a shot after having just made his last two or three shots than he does after having just missed his last two or three shots" (Kahneman and Tversky, 1971). Professional players themselves reported feeling that they "almost can't miss" after making several shots in a row (Kahneman and Tversky, 1971). Players' actions confirm this sentiment, as shot difficulty tends to increase following several made shots (Rao, 2009).

However, are the fans and players really wrong? Is the hot hand fallacy truly fallacious? We believe the question remains an open one. For, buried in the introduction of the famous Gilovich, Vallone, and Tversky paper lies a key assumption (emphasis added):

> "It may seem unreasonable to compare basketball shooting to coin tossing because a player's chances of hitting a basket are not the same on every shot. Lay-ups are easier than 3-point field goals and slam dunks have a higher hit rate than turnaround jumpers. Nevertheless, the simple binomial model is equivalent to a more complicated process with the following characteristics: Each player has an ensemble of shots that vary in difficulty (depending, for example, on the distance from the basket and on defensive pressure), and *each shot is randomly selected from this ensemble*" (Gilovich, Vallone and Tversky, 1985).

The validity of the authors' conclusion hinges heavily on this assumption, and yet it is not difficult to envision a scenario in which it is violated. There is strong evidence that players themselves believe in the hot hand. Thus, players may not select shots at random, but rather choose more difficult shots after they make a few shots because they believe they are hot. Rao shows with a small dataset that this appears to be the case for at least some players (2009). Therefore, the two effects could cancel each other out: players become hot, which increases their shooting percentage overall, but simultaneously choose more difficult shots with a lower probability of going in. On paper, the sign of the effect is ambiguous.

In this paper, we use a novel dataset provided by the SportVU tracking cameras of STATS, Inc., that records the x-y-z coordinates of all players on the court in increments of 1/25th of a second. These cameras, installed in 15 NBA arenas in the 2012-2013 season, provide us with a dataset of over 83,000 shot attempts. Synthesizing this dataset, we are able to know almost all relevant characteristics of the shot at the moment it is taken. This allows us to investigate the following two questions: First, do players (both offensive players and defenders) believe in the hot hand, as evidenced by their playing decisions? And second, if we control for the difficulty of the shot, does the hot hand effect indeed emerge? Although some authors have previously attempted to control for shot difficulty (Rao, 2009), this paper adds to the literature with the sheer scope and quality of the data we have at our disposal. To our knowledge, no analysis of the hot hand has been able to quantify defensive intensity or shot quality to the extent that this dataset allows us to do.

We show that players who perceive themselves to be hot based on previous shot outcomes shoot from significantly further away, face tighter defense, are more likely to take their team's subsequent shot, and take more difficult shots. These results reject the shot selection independence assumption. To account for this bias, we create a comprehensive model of shot difficulty that depends on conditions of the shot the moment it is taken. These conditions include variables relevant to game situation, shot location, and defender locations. Next, we create a measure of heat which reflects the extent to which a player outperformed over his past few shots, bearing in mind how difficult those shots were. By having both a measure of shot difficulty and a measure of heat, we have a way to test for the hot hand, holding shot difficulty constant. The results of this test suggest that once we control for the dependence of shot selection, there may be a small yet significant hot hand effect.

The remainder of this paper is organized as follows: Section 2 presents an in-depth description of this novel dataset, and presents our strategy for controlling for shot difficulty. Section 3 investigates the effect of heat on game play decisions, and Section 4 examines whether the hot hand exists once shot difficulty is controlled for. Finally, Section 5 concludes.

2 The Data

2.1 Raw Data

STATS, Inc. introduced the SportVU optical tracking system into the National Basketball Association ("NBA") in 2010. The system uses six cameras, three on each side of the court, to provide precise three-dimensional image tracking of the players, referees, and ball every 1/25th of a second. During the 2012-2013 season, the system was installed in 15 NBA arenas, fully half of the league.[1] Our dataset consists of every game played at these arenas in the 2012-2013 regular season.

We draw data from four databases: the NBA's Expanded Play-by-Play ("NBA-PBP"), NBA Roster ("Roster"), SportsVU Sequence Optical Tracking ("Optical Tracking") and SportsVU Sequence Play-by-Play Optical (PBP-Optical). All four databases are from the 2012-2013 NBA season and were provided by STATS, although the NBA-PBP data comes from the NBA. We first describe each dataset and then explain how the datasets were combined into our final analysis dataset, the shot log.

NBA Roster

We obtain basic player and team information from the Roster data. This data spans the 2012-2013 NBA season and has information on all 30 teams and 474 players. We extract player traits, including weight, height and position. The dataset also links each player to his respective team.

NBA Expanded Play-by-Play

We obtain data on in-game possessions and events for the 2012-2013 NBA season from the NBA-PBP database. This data is the play-by-play feed that is compiled by the arena's official statistician every game. It lists all major events in the game with additional metadata, such as the time of the

[1] The fifteen teams with SportVU installed are: the Boston Celtics, the Cleveland Cavaliers, the Dallas Mavericks, the Golden State Warriors, the Houston Rockets, the Milwaukee Bucks, the Minnesota Timberwolves, the New York Knicks, the Oklahoma City Thunder, the Orlando Magic, the Philadelphia 76ers, the Phoenix Suns, the San Antonio Spurs, the Toronto Raptors, and the Washington Wizards.

event and the player(s) associated with the event (i.e., the player who took the shot).

The events are tagged with an event-id to indicate the type of event. This allows us to filter the database to only shot attempts. Each shot includes information on the shooter, the shot type, a textual description of the event, the result of the shot, and a series of binary descriptors including whether the shot was a fast-break, blocked, or "in the paint."

SportVU Optical Tracking

We obtain in-game spatial data from the Optical Tracking database. The dataset includes x, y, and z coordinates of each player and the ball at 1/25th of a second increments. The dataset allows us to analyze the movement of each player on both the offense and defense during every possession. We normalize the x and y axes so that coordinates are relative to the basket, which is set at (0,0). This allows us to directly compare action from both baskets. From this data, we can derive information such as shot distance, defender distance, and defender angle.

SportVU Play-by-Play Optical

Each observation of the NBA-PBP is linked to the Optical Tracking data through an observation in the PBP-Optical dataset. This fourth dataset provides the key linkage between the play-by-play feeds, player characteristics, and the optical tracking data. Each shot attempt is given a unique sequence number in both the NBA-PBP and PBP-Optical. We use the precise time stamp (accurate to 1/25th of a second) from the PBP-Optical to link to the Optical Tracking database.

2.2 Shot Log

The analysis shot log synthesizes the four datasets to create robust characterization of each shot. For each shot, we have information from the Roster and PBP-NBA feed on the player who took the shot, the type of shot taken, and the time and score at the time of the shot. From the Optical Tracking Data, we have the precise location of the ball and all ten players, both offensive and defensive, on the court. Table I shows all of the variables included in the shot log.

2.3 Empirical Preliminaries

Before we delve into the full analysis, it is worth laying down some of the empirical groundwork and basic basketball theory that will be used in the subsequent sections of this paper.

2.3.1 Predicted Shot Difficulty

Using the shot log data described in Section 2.2, we estimate a model that predicts the difficulty of each shot for player i taking shot s, based on four broad categories of determinants of shot difficulty:

$$\begin{aligned}\hat{P}_{is} = \alpha + \beta * (Game\ Condition\ Controls_{is}) + \gamma * (Shot\ Controls_{is}) \\ + \delta * (Defensive\ Controls_{is}) + \theta * (Player\ Fixed\ Effects_{i})\end{aligned} \tag{1}$$

We briefly comment on the details and theory behind each of the categories below:

Game Condition Controls

It is plausible that the game situation will affect the difficulty of a shot. Two shots that are identical, save that one was taken in the first quarter with the score tied, and the other taken in the closing seconds of the fourth quarter with the score tied, might have dramatically different difficulties because of differences in pressure or player fatigue. Similarly, identical shots taken at the same time in a game, but with different score differentials may have different difficulties.

As such, we control for game context by including the score differential in the game at the time of the shot, defined as the shooting team's score minus the defending team's score, and dummy variables for each quarter or overtime period in the game. We also include interactions between the score differential and the period dummies.

If players only have a certain amount of effort to give, they may curtail or expend energy based on game situation. In principle, these variables account for variation in shot difficulty that is based on in-game pressure or in-game effort.

Shot Controls

Shot difficulty is clearly a function of distance from the basket and the type of shot that is attempted. In general, the difficulty of a shot increases with distance from the basket. The fuctional form of this relationship, however, is not clear. Factors such as angle from the basket, as well as non-linearities arising from the key or the three-point line, make the relationship between distance and shooting percentage decidedly non-linear.

We account for this by gridding the court into two-by-two foot boxes and include this set of mutually exclusive and exhaustive dummy variables in the shot difficulty model. This allows for a non-parametric specification, where each two foot increment's coefficient is allowed to vary.

As mentioned previously, the NBA-PBP feed provides detailed descriptions of shot type. The shot type data was aggregated into thirteen main categories that are enumerated in Table I. Shot type is a powerful indicator of difficulty above and beyond the location on the floor. To attempt to capture this variation, a set of thirteen mutually exclusive and exhaustive shot type dummies were added to the model. Each of these shot type dummies was interacted with the distance variable to capture variability in the difficulty of a particular shot type at a particular distance.

Finally, we included a binary variable if the shot attempt was labeled as a "Fastbreak" in the NBA-PBP. This generally means that the shooting team was able to force a live-ball turnover from the opposition and quickly transition down the court for a shot attempt. These attempts make up roughly nine percent of the sample, and result in made baskets at a dramatically higher rate - 63 percent vs. 43 percent for non-fastbreak attempts.

Defender Controls

The third class of determinants of shot difficulty covers the role of defensive intensity. The location of defenders, especially the nearest defender, at the time of a shot plays an integral role in determining the likelihood of a shot being made. Using the SportVU data, we are able to determine both the absolute distance between the player shooting and the closest defender and the angle of that defender relative to a straight line between the shooter and the basket (see Figure I).

It is important to note that the SportVU data does not provide us with appendages; the defender distances we use are from the center of the body mass of the player to the center of the body mass of the defender. This is a clear deficiency, as defenders who "close out" on a shooter well with arms upraised to block the shooter's vision of the basket will make a shot more difficult than a defender in the same location who does not have his arm upraised.

We do not constrain the closest defender to be in between the player and the basket. Instead, we attempt to account for the variability in shot difficulty at levels of defender distance by including defender angle and an interaction term between closest defender distance and angle. We also include a binary variable for double coverage, which is defined as the second closest defender being within four feet.

In addition to defender location, we also include data on individual defender-shooter matchup. More specifically, we include the height differential between the closest defender and the shooter. This is interacted with the distance between the two players to give a measure of size mismatch.

Player Fixed Effects

Finally, we use Player Fixed Effects to control for differences between players. If Kevin Durant, one of the best shooters in the NBA, and Tyson Chandler, a center who rarely takes shots from outside of the paint, both take identical jump shots, the two shots likely have different likelihoods of being made because players vary in their shooting abilities. Player Fixed Effects allow us to capture this difference.

Shot Difficulty Regression Results

Table II shows the exact list of controls used for our shot difficulty model. For ease of notation, we will will refer to the estimated probability of a make as simply $\hat{P}$ throughout the paper. We used OLS rather than a Logit or Probit regression because the number of independent variables led to issues with model convergence. Therefore, when the model predicted values of greater than one or less than zero, we replaced them with 0.99 and 0.01 respectively.[2] Because we used such a large number of variables and interactions, along with a very non-parametric specification (i.e. breaking distance into two-by-two zones rather than using a linear measure), most of the hundreds

[2]Note that the predicted probabilities that were outside the bounds of [0,1] represented less than one percent of the sample.

of coefficients do not have an easy interpretation. Thus, no coefficients are reported.[3]

Out of Sample Accuracy Testing

To test the accuracy of the predictions of the model, we split the dataset into a training set (shots taken in the first half of the season) and validation set (shots taken in the second half). We ran our model on the training set, and used our results to predict the $\hat{P}$'s for the validation set, creating a pseudo-out of sample test. We then grouped the $\hat{P}$'s into bins of one percent. For each bin, we calculated the percentage of shots that were actually made. If the model is accurate, the $\hat{P}$ bins should correspond closely with the actual make percentages. Figure II presents the scatter plot of the data. Aside from some under-prediction for very low values of $\hat{P}$, the model fits the data very well out of sample.

$\hat{P}$ gives us a single number that is both easily interpretable, and encapsulates the overall "difficulty" of a given shot. This will make our subsequent analyses simpler and easier to understand.

It is worth noting that if the hot hand truly does exist, there is potential for bias in these $\hat{P}$ values. If players attempt more difficult shots when they are hot and are more likely to make these shots when they are hot, then our model may overestimate the probability of difficult shots and underestimate the probability of easy shots as taken by a "neutral player." However, if this bias does exist, it would make the difficulty-controlled hot hand *more* difficult to detect. In other words, if we find an effect, it is in spite of this bias, not as a result of it.

2.3.2 Defining Heat

Before considering the effect of heat, we must first define what heat actually means. Note that we only consider consecutive shots that occur within the same game. The conventional definition of heat, which we refer to as Simple Heat, is as follows:

$$Simple\ Heat_n = \text{ Actual \% over past } n \text{ shots} \tag{2}$$

[3]There is a small worry that using the entire sample might bias the $\hat{P}$'s because the individual shot outcome appears on both sides of the regression equation, but because we have over 80,000 shots, this effect should be *de minimis*. The correlation between the $\hat{P}$'s estimated on the full sample and those estimated on half the sample in the out of sample testing below is 0.973.

However, we believe that in order to test for the hot hand, we must define what we call Complex Heat. Note that the $\hat{P}$ values allow us to calculate the expected shooting percentage over the past n shots. Therefore, we define:

$$Complex\ Heat_n = \text{ Actual \% over past } n \text{ shots} - \text{ Expected \% over past } n \text{ shots} \tag{3}$$

Simple Heat is the commonly understood measure of heat - it simply reflects how successful a player has been over his past few shots. However, a drawback of this is that it does not account for shot difficulty. The implications of this drawback are illustrated by an example outlined in Figure III: Simple Heat rates Thad Young as hot because he has made each of his last five layups and dunks, but based on the $\hat{P}$ of each of those shots, we would expect him to make 87% of those shots. By contrast, Klay Thompson only makes three of five very difficult shots, but he has exceeded his expectation by more than Young. Complex Heat thus allows us to look for players who are shooting better than expected, given the difficulty of the shots they are taking.

Another drawback of Simple Heat is illustrated by the following example. Suppose a player is playing against a very short defender. It is likely that he is making more of his shots than usual, and his Simple Heat will be high. However, as long as he keeps shooting against this low-quality defender, his shots are higher-probability of going in. In other words, $\hat{P}$ will also be high. Therefore, we see that Simple Heat and $\hat{P}$ have a mechanical correlation, which biases the estimated effect of heat down. This is not the case with Complex Heat, because it controls for the difficulty of the past few shots (i.e. the $\hat{P}$ of the previous few shots). Continuing with the short defender example, the defender's short stature will be controlled for when we use Complex Heat, because the difficulty of the past few shots will be lower.

If $\hat{P}$ is specified correctly, the Complex Heat specification can essentially be thought of as controlling for factors that are relatively constant across stretches of games. These factors could include defender matchups, offensive or defensive game plans, effort exerted (i.e., blowouts or close games), and fatigue. This idea will be explored further in the following section.

Finally, we note that there is some discretion in selecting how many shots we "look back" over when defining heat (i.e. what value n takes when defining Simple and Complex Heat). We ran our results for all values of n from two through seven, and found similar results. For simplicity, we report all results here for $n = 4$, which we believe is a reasonable number of shots. Moreover, unless otherwise noted, all results that use heat refer to Complex Heat, which we believe is the more correct measure.

3 Do Players Believe in the Hot Hand?

The first question we are interested in is: Do players believe in the hot hand? Gilovich, Vallone and Tversky (1985) showed that players on the 76ers claimed to believe the hot hand was true. We are interested in whether over a quarter century later, players' actions reflect this belief.

3.1 Empirical Strategy

To understand players' responses to heat, we look at the effect of heat on shot distance, closest defender distance, and overall shot selection. Given that Simple Heat is how players and coaches most likely conceive of the hot hand, we test both Simple and Complex Heat. The results are consistent across both measures of heat.

3.1.1 Shot Distance

If players truly believe that the hot hand exists, they may attempt more difficult shots as they heat up. One way to test this is to see whether players take shots that are further away from the basket as they become hot. To see if this is the case, we run the following specification:

$$Shot\ Distance_{is} = \alpha + \beta * (Heat_{is}) + \gamma * (Controls_{is}) + \theta * (Player\ Fixed\ Effects_i) \quad (4)$$

The controls include quarter, score differential, quarter/score differential interaction, closest defender distance, and fast break. Shot type is notably missing - this is because shot type and shot distance are highly related, and so it doesn't make sense to have both variables on opposite sides of the equation.

We would hypothesize that if players believe in the hot hand, the coefficient on heat would be positive. The intuition is that as players think they are becoming hot (i.e. better shooters), they optimize by taking shots that are more difficult than shots they would ordinarily attempt. However, an alternative explanation for a positive coefficient is that the defense buys into the hot hand belief. Therefore, they cover the player more tightly, and don't allow them to take any shots from close range. It is difficult with this specification to distinguish exactly which set of players are reacting to perceived heat.

3.1.2 Defender Distance

As mentioned previously, if defenders believe the hot hand exists, they may cover hot players more tightly. To test this hypothesis, we run the following regression:

$$Defender\ Distance_{is} = \alpha + \beta * (Heat_{is}) + \gamma * (Controls_{is}) + \theta * (Player\ Fixed\ Effects_i) \quad (5)$$

The controls in this regression include quarter, score differential, quarter/score differential interaction, fast break, shot type, and shot distance.

We would hypothesize that if defenders believe in the hot hand, then the coefficient on heat would be negative. Intuitively, as a player becomes hotter, the defenders give him less space to shoot. Alternatively, the hot player is willing to take more closely guarded shots as he heats up.

3.1.3 Likelihood of Taking Next Shot

Further, we can investigate the effect of heat on overall shot selection by players by evaluating how heat impacts the probability that a given player takes his team's next shot. More specifically, we look at the probability that after Player X takes a shot, Player X also takes his team's next shot as a function of heat. To do this, we estimate the following Probit model:

$$P(Same_{is}) = \Phi(\alpha + \beta * (Heat_{is}) + \gamma * (Controls_{is}) + \theta * (Player\ Fixed\ Effects_i)) \quad (6)$$

In this regression, we control for "game environment" factors - quarter, score differential, and the quarter/score differential interaction. Here, if players believe in the hot hand, we would expect the coefficient on heat to be positive. If the shooter thinks he is heating up, he is more likely to shoot. Moreover, if his teammates also believe he is hot, they are more likely to give him the ball in a shooting position.

3.1.4 Overall Shot Difficulty

Finally, we can use our $\hat{P}$ model to investigate how heat effects overall shot difficulty. If players attempt shots that are further away and more closely guarded when they perceive themselves to be hot, we would expect that these shots overall have a lower value of $\hat{P}$. To see if this is the case, we run:

$$\hat{P}_{is} = \alpha + \beta * (Heat_{is}) \tag{7}$$

Note that we don't include controls or player fixed effects, since $\hat{P}$ includes these already. If players do indeed attempt all-around more difficult shots when hot, we would expect β to be negative (recall that a low $\hat{P}$ corresponds to a more difficult shot).

3.2 Results

3.2.1 Shot Distance

Looking at Table III, we see that shot distance increases with heat. This effect appears robust regardless of whether we use Simple or Complex Heat. The size of the effect is significant - one way to think about this is to consider the coefficient of 2.293 on Simple Heat in column (1). If a player makes one more of his past four shots (i.e. increases Simple Heat by 0.25), his estimated shot distance increases by about 0.57 feet, or nearly seven inches. Given that the average shot in the sample is 12.7 feet from the basket, this represents a 4.5% increase. The effect size is consistent for Complex Heat as well.

This supports the hypothesis that as players become hot, they attempt more difficult shots. However, it is unclear exactly what the source of this effect is. Perhaps the most obvious explanation is that shooters believe in the hot hand, become more confident, and attempt shots that are further away from the basket than they would ordinarily take. However, another compelling explanation is that the defense also buys into the hot hand, and tightens their coverage of the hot player. As a result, the player is only permitted to take shots that are further from the basket.

3.2.2 Defender Distance

Considering Table III, we also see that defender distance shrinks with heat. Once again, this effect is consistent across both specifications of heat. Again, to better understand the size of these coefficients, consider the coefficient on Simple Heat in column (2). If a player makes one more of his past four shots, estimated defender distance shrinks by 0.053 feet (over half an inch). This is a small effect, but it is worth noting that the average defender distance is only about four feet. Therefore, the effect size is about equal to a 1% decrease in defender distance.

This supports the hypothesis that as a player becomes hot, the defense plays tighter defense on that player. However, as discussed above, it is difficult to disentangle the defensive and offensive response. A less obvious but equally valid explanation is that hot players simply attempt shots with tighter coverage that they would ordinarily pass up. Regardless, this supports that players' actions are consistent with a belief in the hot hand.

3.2.3 Overall Shot Selection

First we consider how likely it is that a player takes his team's next shot. Table III shows how the probability that the same player takes a team's next shot as a function of heat. Regardless of how heat is specified, the marginal effect is between 0.055 and 0.061. Again, using Simple Heat, if a player makes one more of his past four shots, his probability of taking his team's next shot increases by 1.4 percentage points, assuming all covariates are at their mean values. This sounds small, but remember that there are five players on a team. If the average player has about a 20% chance of taking his team's next shot, so this corresponds to a 7% increase in the overall probability.

The effect of heat on overall shot difficulty is initially less clear. Looking at Table III, it appears that when we use Simple Heat, heat is associated with taking an easier (higher-probability shot). This is the opposite of what we would expect if players believe in the hot hand and react accordingly. However, when we use Complex Heat, the coefficient flips sign.

Although at first confusing, this matches our discussion of the bias associated with using Simple Heat. Returning to the short defender example, we can explain this counterintuitive result: Suppose a player is being guarded by a very short player. He will most likely be making lots of shots, and have a high value of Simple Heat. He may react to this by taking his next shot from further away, and his defender may guard him closer (as evidenced by Tables III and III), but he will still be covered by this short defender. As a result, this next shot will still be relatively easy. In other words, we will see a positive relationship between Simple Heat and $\hat{P}$. This won't be an issue with Complex Heat, however, because we control for the difficulty of these past shots. This is evidenced by the negative coefficient on Complex Heat.

To confirm that this is the correct explanation, we decompose Complex Heat into its respective pieces: Actual shooting percentage (i.e. Simple Heat) and expected shooting percentage. Table IV shows these results. Here we see that expected shooting percentage has a positive coefficient, while shooting past percentage (Simple Heat) now has a negative coefficient. This is most likely explained by serial correlation between shot difficulty, as induced by factors like the short defender. However, once we control for past shot difficulty, making more shots (becoming hot) is correlated with taking more difficult shots, as we expected.

4 Testing the Hot Hand

Now armed with evidence that shot difficulty increases with heat, we can turn our attention to whether the hot hand truly exists once we control for shot difficulty. In the following section, we present results using our $\hat{P}$ shot difficulty model, comment on those results, and add some discussion of potential drawbacks of this approach.

4.1 Empirical Strategy

First, as a baseline, we look for the hot hand without any control for shot difficulty. We are interested in how a player's probability of hitting a shot varies with heat and nothing else. Therefore, we can run the following simple specification:

$$P(Make_{is}) = \alpha + \beta * (Heat_{is}) + \theta * (Player\ Fixed\ Effects_i) \tag{8}$$

We can run this using an OLS specification to understand how the probability of making a shot varies with heat, unconditional on shot type or difficulty. This is analogous to the analysis done by Gilovich, Vallone, and Tversky (1985) and others, and will provide us with a baseline.

Next, we can control for difficulty by using $\hat{P}$, which encapsulates all of the relevant and quantifiable controls. We can use the following relatively simple specification to test for the difficulty-controlled hot hand:

$$P(Make_{is}) = \alpha + \beta * (Heat_{is}) + \gamma * \hat{P}_{is} \tag{9}$$

Note that we no longer need player fixed effects, because they are contained within $\hat{P}$.

If the hot hand does not exist, and each shot is truly independent, we would expect $\alpha = 0$, $\beta = 0$, and $\gamma = 1$. In other words, the only thing that predicts $P(Make)$ is the difficulty of the shot. However, if the hot hand does exist, we would expect to find $\beta > 0$.

4.2 Results

In Table V, column (1) loosely replicates the work done by Gilovich, Vallone, and Tversky by regressing a simple measure of heat and player fixed effects against the probability of hitting a shot, with no attempt to control for shot difficulty. Our results mirror the original authors', with the coefficient on heat being negative but insignificant. Column (2) repeats this analysis but uses Complex Heat, but the results remain similar.

Column (3) introduces shot difficulty controls via $\hat{P}$, but still uses Simple Heat rather than Complex Heat. As discussed previously, Simple Heat alone may be a misleading measure of heat because it does not account for factors that affect shot difficulty and remain constant across game stretches. In column (3), we actually see a supposed significant "cold hand" effect. This dramatic shift is evidence of the bias of Simple Heat: it does not control for correlation across series of shots, such as the short defender example. Therefore, we believe that Complex Heat is a superior measure for testing for the hot hand.

As shown in the previous section, shot difficulty increases as a player "heats up." Therefore, it seems possible that once we control for difficulty, the hot hand effect would indeed emerge. When we control for difficulty and use Complex Heat, we see a positive and significant hot hand effect. We acknowledge that the effect size is modest. The coefficient of 0.0211 means that a player who makes one more of his past four shots sees his shooting percentage increase by 0.53 percentage points. Given that the average NBA player has a field goal percentage of about 45%, this represents about a 1.2% improvement. In the same vein, if a player makes two more of his past four shots (perhaps more indicative of what it truly means to be "hot"), we see a 2.4% improvement.

While the absolute size of the measured effect is small, we have reason to believe the true effect could be larger. As discussed by Daniel Stone (2012), if shot outcome (shooting percentage) is a noisy measure of true underlying heat, then error-in-variables bias will tend to push our heat coefficient closer to zero.

4.3 Discussion

Based on these results, concluding that the hot hand exists is contingent on defining "hot" correctly. A player who makes three out of his past four layups is not hot, but a player who makes three out of his past four three-point attempts is. In other words, being hot is not about the *absolute* number of shots a player has previously made, but rather is about how much he has outperformed, conditional on the types of shots he has taken.

There are several drawbacks to the two stage empirical strategy we employ here of first estimating a shot difficulty ($\hat{P}$) regression and then using those results to test for the hot hand. The principal worry is errors-in-variables: if our $\hat{P}$'s are measured too imprecisely, they may produce a biased estimator. There is reason to suspect that, despite our best efforts and the extraordinary dataset at our disposal, our shot difficulty model does not control for or does not correctly specify individual shot difficulties. Specific concerns include: not being able to track appendages, not specifying relationships between variables (i.e., distance from the basket, defender distance, etc.) correctly, and that player fixed effects are not precise enough to accurately estimate $\hat{P}$ for individual players. Player fixed effects change the intercept of the probability that a shot goes in for each player, but they do not account for the fact that certain players "specialize" in certain shots - in other words, they do not adjust the slopes.

Additionally, we must acknowledge that the formulation of the hot hand supported here is not what is typically associated with the hot hand phenomenon. It is absolute outperformance that typically comes to mind when discussing the hot hand rather than this relative outperformance. Therefore, we are left to conclude that the hot hand appears to exist in some form, but it may not be the effect most people envision when they talk about the hot hand.

5 Conclusion

For thirty years, the empirical consensus that the "hot hand" in basketball is a fallacy of the human mind has been confirmed time and time again. In the same way that evolutionary biologists might regard creationists as completely misguided, economists, psychologists and statisticians have viewed the persistent belief in the hot hand as utterly fallacious. Amos Tversky, co-author of the canonical paper on the subject, typifies this view when he says, "I've been in a thousand arguments over this topic, won them all, but convinced no one" (Bar-Eli, Avugos and Raab, 2006). For the first time, however, thanks to SportVU's optical tracking data, we have the dataset necessary to challenge the central assumption of shot selection independence that underlies most of this literature.

In this paper, we show that the shot selection independence assumption is not a good one. Players who have either simply made more of their last few shots, or done better than expected on those shots, tend to shoot from farther away and with the nearest defender closer. These shots are consequently more difficult. Additionally, hot players are much more likely to take the team's next shot, indicating that players are taking more opportunities to shoot when they believe that they are hot, and thus are not choosing shots independently.

We then extend our analysis to ask if the hot hand exists once we control for this dependent shot selection using a regression framework that controls for past expectation. This approach finds a small, positive, and significant hot hand effect. This conception of the hot hand as exceeding expectations is different from the popular conception of absolute outperformance. Given the endogeneity of shot selection, however, assessing the hot hand only through the absolute outperformance seems flawed.

We note that this novel empirical strategy has potential applications outside of simply evaluating the hot hand. They could be used at the player level to create all-inclusive metrics of individual offensive (or defensive) ability. Shot difficulty provides a far more precise measures of shooting ability than current field goal percentage statistics do.

We caution that our use of OLS regression as the functional form of choice for most of our analysis may not be optimal. Moreover, our effect sizes are small, and it is more likely that our estimates of standard errors are imprecise than our estimates of mean effects. It is plausible that better specifications might render these small affects insignificant. Nevertheless, our effects remain significant across different specifications and with both clustered and heteroskedasticity robust standard errors. Future extensions might use more complicated functional forms or modeling techniques.

We hope this paper will spark debate and further work on the mechanisms of the hot hand effect in basketball. Understanding how both players and defenses react to perceived hotness can provide valuable insight into teams' and players' optimal strategies and could lead to a re-evaluation of current strategies. At the very least, our findings cast doubt on the overwhelming consensus that the hot hand is a fallacy. Perhaps the next time a professor addresses the Harvard Men's basketball team, the hot hand will not be so quickly dismissed.

References

[1] Adams, Robert M. "The 'Hot Hand' Revisited: Successful Basketball Shooting as a Function of Intershot Interval." *Perceptual and Motor Skills,* 74 (1992), 934.

[2] Arkes, Jeremy. "Revisiting the Hot Hand with Free Throw Data in a Multivariate Framework." *Journal of Quantitative Analysis of Sports,* 6 (2010), 1-12.

[3] Bar-Eli, Michael, Simcha Avugos, and Markus Raab. "Twenty Years of 'Hot hand' Research: Review and Critique." *Psychology of Sport and Exercise,* 7 (2006), 525-553.

[3] Camerer, Colin F. "Does the Basketball Market Believe in the Hot Hand?" *American Economic Review,* 79 (1989), 1257-1261.

[4] Croson, Rachel, and James Sundali. "The Gambler's Fallacy and the Hot Hand: Empirical Data from Casinos." *Journal of Risk and Uncertainty,* 30 (2005), 195-209.

[5] Davidson, Adam. "Boom, Bust or What? Larry Summers and Glenn Hubbard Square Off on Our Economic Future." *New York Times*, May 2, 2013.

[6] Gilovich, Thomas, Robert Vallone, and Amos Tversky. "The Hot Hand in Basketball: On the Misperception of Random Sequences." *Cognitive Psychology,* 17 (1985), 295-314.

[7] Green, Brett S. and Jeffrey Zwiebel. "The Hot Hand Fallacy: Cognitive Mistakes or Equilibrium Adjustments? Evidence from Baseball." Stanford University Graduate School of Business Research Working Paper, (2013).

[8] Huizinga, John and Sandy Weil. "Hot Hand or Hot Head? The Truth About Heat Checks in the NBA." MIT Sloan Sports Analytics Conference (2009).

[9] Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser. "Hot Hands in Mutual Funds: Short-Run Persistence of Relative Performance, 1974-1988." *Journal of Finance,* 48 (1993), 93-130.

[10] Kahneman, Daniel, and Amos Tversky. "Subjective Probability: A Judgment of Representativeness." *Cognitive Psychology*, 3 (1972), 430-454.

[11] Kahneman, Daniel, Paul Slovic, and Amos Tversky. *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge: Cambridge University Press, 1982).

[12] Koehler, Jonathan, and Caryn Conley. "The 'Hot Hand' Myth in Professional Basketball." *Journal of Sport & Exercise Psychology,* 25 (2003), 253.

[13] Rao, Justin M. "Experts' Perceptions of Autocorrelation: The Hot Hand Fallacy Among Professional Basketball Players." 2009.

[14] Stone, Daniel F. "Measurement Error and the Hot Hand." *The American Statistician,* 66 (2012).

[15] Yaari, Gur, and Shmuel Eisenmann. "The Hot (Invisible?) Hand: Can Time Sequence Patterns of Success/Failure in Sports be Modeled as Repeated Random Independent Trials?" PloS One, 6 (2011).

Tables

Table I: Shot Log Covariates

Variable Category	Variable	Description
Game Conditions	Quarter	Dummy for current quarter (or overtime period)
	Score Differential	What is the current score differential - is the game close?
	Time Remaining	The time, in minutes and seconds, left on the clock at the time of any event
Shot Conditions	Shot Location	Both a series of dummies for each 1x1 foot square on the court, and a precise to three decimal places (x,y) coordinate location
	Fast Break?	Dummy for fast break (i.e. undefended) shot
	Shot Type	Series of dummies - 13 (mutually exclusive) categories:
		Dunk
		Tip-in
		Driving layup
		Reverse layup
		Layup
		Hook shot
		Bank shot
		Turnaround
		Pull-up
		Floater
		Fadeaway
		Stepback
		Jumpshot
Defensive Conditions	Defender Distance	Absolute distance between shooter and first, second, third, fourth, and fifth closest defenders
	Defender Angle	Angle between shooter's direct line to the basket and the closest defender
	Height Difference	Height difference between shooter and closest defender (proxy for defensive mismatch)
	Double Covered?	Is the second-closest defender within 4 feet of the shooter?

Table II: Predicting Makes - List of Variables Used in Specification

VARIABLES	
Quarter Dummies	
Score Differential	
Shot Location Dummies	
Shot Category Dummies	
Fastbreak Dummy	
Distance of Closest Defender	
Angle of Closest Defender	
Shooter-Defender Height Difference	
Double Covered Dummy	
Quarter × Score Differential	
Shot Location Dummies × Shot Category Dummies	
Distance of Closest Defender × Angle of Closest Defender	
Distance of Closest Defender × Shooter-Defender Height Difference	
Player Fixed Effects	
Observations	70,862
R^2	0.149

Table III: Do Players Believe in the Hot Hand?

(a) Simple Heat

VARIABLES	(1) Shot Distance	(2) Defender Distance	(3) P(Same)	(4) $\hat{P}$
Simple Heat	2.293***	-0.213***	0.0551***	0.0509***
	(0.181)	(0.0392)	(0.00888)	(0.00358)
Constant	6.969***	4.226***		0.433***
	(0.286)	(0.140)		(0.00191)
Observations	43,475	43,475	43,467	43,475
R^2	0.296	0.167		0.005

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

(b) Complex Heat

VARIABLES	(1) Shot Distance	(2) Defender Distance	(3) P(Same)	(4) $\hat{P}$
Complex Heat	2.147***	-0.184***	0.0613***	-0.0418***
	(0.180)	(0.0392)	(0.00927)	(0.00380)
Constant	8.067***	4.126***		0.457***
	(0.262)	(0.138)		(0.000904)
Observations	43,475	43,475	43,467	43,475
R^2	0.296	0.167		0.003

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Simple and Complex heat are calculated over the past four shots. Shot distance regressions include controls for quarter, score differential, quarter/score differential interaction, closest defender, and fast break. Defender distance regressions include controls for quarter, score differential, quarter/score differential interaction, fast break, shot type, and shot distance. Likelihood of taking next shot regressions include controls for quarter, score differential, and quarter/score differential interacton. Shot distance, defender distance, and likelihood of taking next shot regressions all have standard errors clustered by player.

Table IV: Decomposing Complex Heat

VARIABLES	(1) P(Make)
$\hat{P}$	1.051***
	(0.0101)
Actual Shooting Percentage	0.0134
	(0.00948)
Expected Shooting Percentage	-0.221***
	(0.0209)
Constant	0.0699***
	(0.00928)
Observations	43,475
R^2	0.151

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Actual and expected shooting percentages are calculated over the past four shots.

Table V: Does the Hot Hand Exist?

VARIABLES	(1) P(Make)	(2) P(Make)	(3) P(Make)	(4) P(Make)
Simple Heat	-0.0110		-0.0240***	
	(0.0108)		(0.00875)	
Complex Heat		-0.0190*		0.0211**
		(0.0113)		(0.00947)
$\hat{P}$			1.020***	1.020***
			(0.00984)	(0.00982)
Constant	0.459***	0.454***	-0.000465	-0.0115**
	(0.00507)	(8.60e-05)	(0.00629)	(0.00505)
Observations	43,475	43,475	43,475	43,475
R^2	0.014	0.014	0.149	0.149

Robust standard errors in parentheses

Notes: Simple and Complex heat are calculated over the past four shots. Regressions that do not include $\hat{P}$ (columns (1) and (2)) include player fixed effects and have standard errors clustered by player.

Figures

Figure I: Illustrating Defender Angle

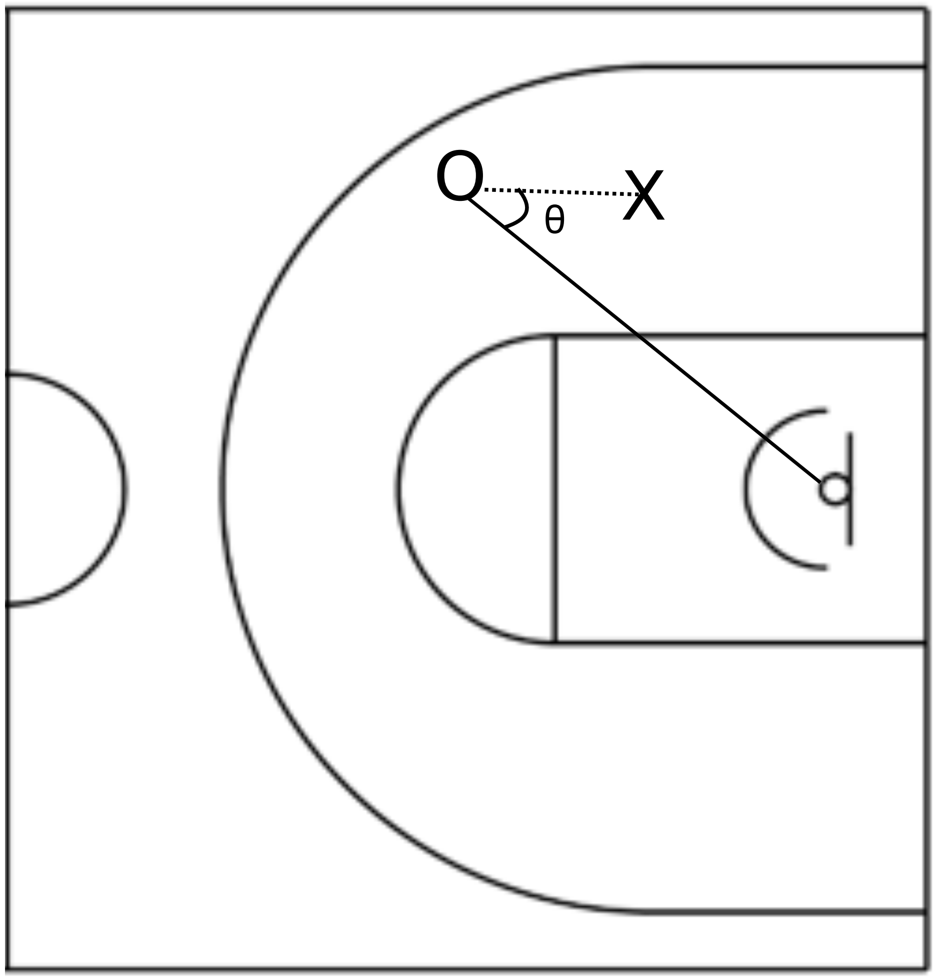

Notes: The O represents the shooter, while the X is the defender. The solid line is the distance to the basket and the dotted line is the distance between the shooter and the defender. θ represents the angle between the defender and the shooter's direct line to the basket.

Figure II: Predicted vs. Actual Shooting Percentages

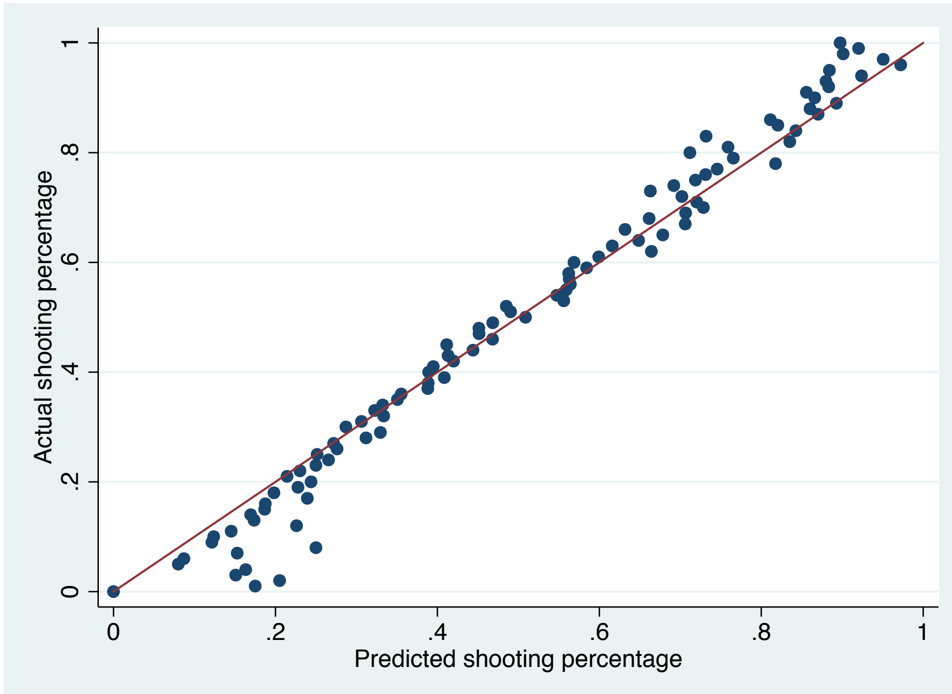

Notes: The x-axis is the estimated shooting percentage of shots taken in the second half of the season. These values are generated by running the model on the first half of the dataset (shots taken in the first three months of the season) and then applying it to the second half (shots taken in the last three months of the season). The y-axis is the actual shooting percentage of these shots. The line y=x represents a perfect fit; we see the dots fall fairly close to the line, indicating our model is indeed predictive.

Figure III: Simple versus Complex Heat

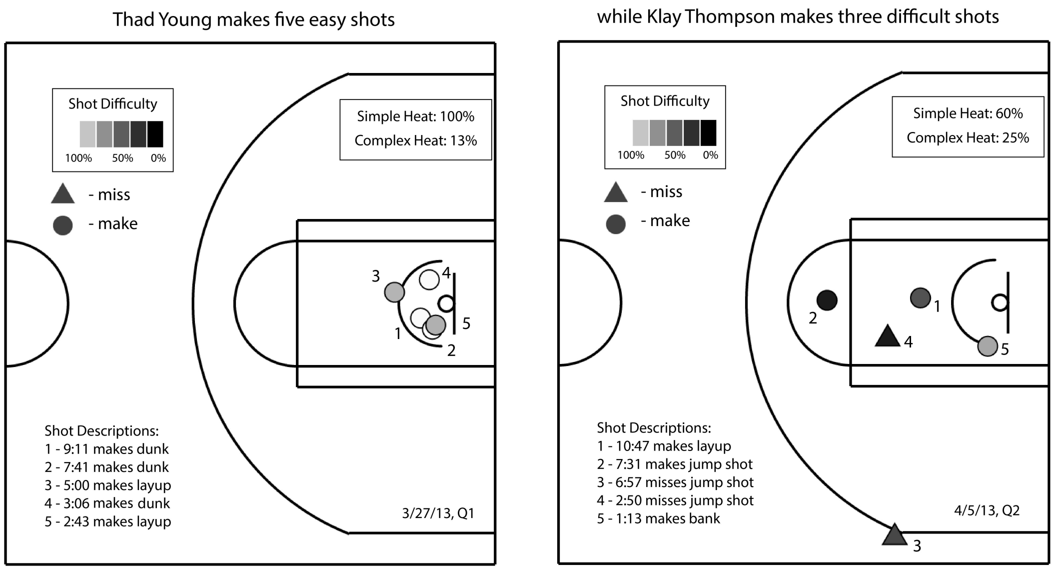

Notes: On the left, we see Thad Young make five easy shots in a row. Since he went five for five, his Simple Heat value is 100%. However, since these shots were not difficult, his expected shooting percentage was 87%, and so his Complex Heat is just 13% (100% - 87%). On the right, we see Klay Thompson make three difficult shots. Since he only went three for five, his Simple Heat is just 60%. However, given the difficulty of these shots, we only expected him to shoot 35%. Thus his Complex Heat is 25% (60% - 35%). This example illustrates how Complex Heat gives players "credit" for how hard their shots are. A player can have a lower overall shooting percentage but still be "hotter" as measured by Complex Heat, as illustrated here.