

Basketball efficiency metrics

Michael Lopez, Skidmore College

Overview

In this lab, we'll gain experience implementing traditional and more novel approaches to measuring shooter accuracy. We'll also implement newer approach for looking at the repeatability of a statistic. First, our preamble to get shot level data from the 2014-2015 NBA season.

```
library(RCurl)
library(mosaic)
url <- getURL("https://raw.githubusercontent.com/JunWorks/NBAstat/master/shot.csv")
nba.shot <- read.csv(text = url)
nba.shot <- na.omit(nba.shot)
nba.shot <- filter(nba.shot, PTS <4, SHOT_DIST>=22 | PTS_TYPE==2, SHOT_DIST <=22 | PTS_TYPE ==3)
nrow(nba.shot)
```

Let's get player specific rate statistics. These include field goal percentage (FGP) and effective field goal percentage (eFGP). (Note: I don't have free throw info, so we can't look at true shooting percentage).

Additionally, `n.shots` stores the number of shots for each player. We'll filter so that we are only looking at players with at least 100 shots.

```
player.shoot <- nba.shot %>%
  group_by(playerName) %>%
  summarize(n.shots = length(FGM),
            FGP = sum(FGM)/(n.shots),
            eFGP = sum(PTS)/(2*n.shots))
```

Let's look at some of our top shooters and our worst shooters, as judged by eFGP.

```
head(arrange(player.shoot, eFGP))
head(arrange(player.shoot, -eFGP))
```

Before we go much further, it's pretty clear there's an issue with players who took a small number of shots. We can also visualize this using a scatter plot.

```
xyplot(eFGP ~ n.shots, data = player.shoot)
```

1. Why does the scatter plot fan in? Is that what we expected?

Let's restrict our sample to only shooters with at least 500 shots.

```
player.shoot <- filter(player.shoot, n.shots >=500)
xyplot(FGP ~ n.shots, data = player.shoot)
xyplot(eFGP ~ n.shots, data = player.shoot)
```

2. Does there appear to be a link between FGP and number of shots or eFGP and number of shots?

Next, we look at the link between eFGP and FGP.

```
xyplot(eFGP ~ FGP, data = player.shoot)
```

3. Why is there a line in the scatter plot above at the line $y = x$? What players are reflected in those points?

Repeatability of shooting percentage metrics.

With baseball and football metrics, we frequently looked at how well a certain statistic correlated with itself (or another statistic) in the future. This is often useful for decision makers in sports who are looking to predict a player's future given his or her past.

In the NBA data set, however, we *don't* have multiple years of data. Instead, we'll split our larger data frame of shots into two smaller data frames, each with half of the shots, and compare the associations between shooting metrics between the two smaller subsets. In essence, this looks at the repeatability of a metric within the same time frame, in place of looking forward.

First, we split the larger data set into two smaller ones, using a random number generator and the filter command. In this example, `nba.shot1` stores the first subset, while `nba.shot2` stores the second.

```
set.seed(1)
nba.shot$test <- rnorm(nrow(nba.shot))>0
nba.shot1 <- filter(nba.shot, test)
nba.shot2 <- filter(nba.shot, !test)
nrow(nba.shot1)
nrow(nba.shot2)
```

There are roughly 160 more shots in `nba.shot1`, but that's fine, and just due to chance.

Next, we repeat the same procedure as above within each the two smaller data frames.

```
player.shoot1 <- nba.shot1 %>%
  group_by(playerName) %>%
  summarize(n.shots1 = length(FGM),
            FGP1 = sum(FGM)/(n.shots1),
            eFGP1 = sum(PTS)/(2*n.shots1))

player.shoot2 <- nba.shot2 %>%
  group_by(playerName) %>%
  summarize(n.shots2 = length(FGM),
            FGP2 = sum(FGM)/(n.shots2),
            eFGP2 = sum(PTS)/(2*n.shots2))

head(player.shoot1)
head(player.shoot2)
```

4. How many of A.J. Price's shots ended up in the first subset of shots? And how many in the second?

Finally, we join the two player-level shot subsets.

```
player.shoot.all <- inner_join(player.shoot1, player.shoot2)
player.shoot.all <- filter(player.shoot.all, n.shots1 >=100, n.shots2 >=100)
```

5. Identify the correlations between (FGP1 and FGP2) and (eFGP1 and eFGP2). Which is more repeatable: a player's shooting percentage or his effective field goal percentage?
6. Provide a logical explanation for why eFGP may be less repeatable than field goal percentage.
7. Why might effective field goal percentage still be a preferred metric, even if it is less repeatable?

Pretty graphs

```
fit.1 <- glm(SHOT_RESULT == "made" ~ SHOT_DIST + TOUCH_TIME +
            DRIBBLES + SHOT_CLOCK + CLOSE_DEF_DIST,
            data = nba.shot, family = "binomial")

nba.shot <- nba.shot %>%
  mutate(predicted.probs = fitted(fit.1),
         expected.pts = predicted.probs * PTS_TYPE,
         epa = PTS - expected.pts)

shot.group <- nba.shot %>%
  group_by(playerName) %>%
  summarise(total.epa = sum(epa)) %>%
  arrange(total.epa)

all.shooters <- inner_join(shot.group, player.shoot)

p <- ggplot(all.shooters, aes(FGP, total.epa, label = playerName))
p + geom_text() + scale_y_continuous("Expected Points Added") +
  scale_x_continuous("eFGP") +
  ggtitle("Expected points added ~ effective FGP, 2014-15 season") +
  theme_bw()

p <- ggplot(all.shooters, aes(FGP, total.epa, label = playerName))
p + geom_text() + scale_y_continuous("Expected Points Added") +
  scale_x_continuous("eFGP") +
  ggtitle("Expected points added ~ effective FGP, 2014-15 season") +
  theme_bw()
```

8. Is the association between expected points added (EPA) stronger when compared to field goal percentage or effective field goal percentage. Why does this follow our intuition?