Journal of Economic Behavior & Organization 136 (2017) 107–124

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Optimal stopping in the NBA: Sequential search and the shot clock[☆]

Matt Goldman[a], Justin M. Rao[a,b,*]

[a] Microsoft AI and Research, United States
[b] U. Washington, United States

ARTICLE INFO

Article history:
Received 5 December 2016
Received in revised form 10 February 2017
Accepted 17 February 2017
Available online 27 February 2017

JEL classification:
D03
C57
D83

Keywords:
Optimal stopping
Expertise
Lab vs. field
Behavioral biases
Marginal thinking

ABSTRACT

We study how experienced agents solve a sequential search problem. In professional basketball teams must shoot within 24 s of the start of a “possession.” The decision of when to shoot requires weighing the current shooting opportunity against the continuation value of a possession. At each second of the “shot clock,” optimal play requires that a lineup's reservation shot value equals the continuation value. We empirically test this prediction with a structural stopping model. Most lineups adopt a reservation threshold that matches the continuation value closely. Overall, the lineups we study capture 84% of the gains of a dynamic vs. an optimal fixed threshold. Lineups with more shared playing experience performed better on average. Observed mistakes lean towards “impatience” – the adopted threshold is either in too low or has excess steepness – meanings too many shots are taken early in the possession.

1. Introduction

In the National Basketball Association (NBA), each time a team possesses the ball they must shoot within 24 s. Failing to do so results in zero points and a loss of the ball. This rule embeds a sequential search problem at the core of game-play. Economic theory posits that the offense should use a “reservation threshold,” based on the time remaining on the shot clock, that determines if a “scoring opportunity” should be used or the possession should be continued in the hopes of observing a better opportunity prior to the expiration of the “shot clock.” The logic of this choice, comparing a current opportunity to the option value of continuing a dynamic interaction, underlies the theoretical requirements of rational decision making in many contexts, such as repeated games and bargaining (Rust, 1987; Provencher, 1997; Arlotto et al., 2013). Laboratory subjects tend to struggle with these types of tradeoffs (Ochs and Roth, 1989; Lee, 2006; Binmore et al., 2002), raising the question of whether people can adhere to the predictions of economic theory.

NBA players are paid millions of dollars a year, face the “shot clock problem” thousands of times per season and must make quick decisions without the help of mathematical software. They thus form an attractive population on which to investigate

☆ We would like to thank James Andreoni, David Eil, Pedro Rey Biel and Charlie Sprenger for extensive notes. Nageeb Ali, Gordon Dahl, Uri Gneezy, Craig McKenzie, David Miller, William Peterman, Joel Sobel and Joel Watson provided helpful comments..

* Corresponding author at: U. Washington, United States.
E-mail addresses: mattgold@microsoft.com (M. Goldman), justin.rao@microsoft.com (J.M. Rao).

http://dx.doi.org/10.1016/j.jebo.2017.02.012

this question. We do so using gameplay data from 2006–2010. The data include the players on the court, the timing and location of shots and other key actions. We model the behavior of an offensive "lineup" (teammates on the court) facing the choice to either "use" the current scoring opportunity or continue the possession.[1] A lineup's *reservation threshold* specifies, for each point of the shot clock, the lowest expected point value a lineup is willing to accept as compared to continuing the possession. This threshold represents a team's marginal shot, because to shoot slightly more frequently they would have to accept opportunities equal to the threshold in expected value. Under straightforward conditions, we show that optimal play requires the reservation threshold is set equal to the *continuation value* of the possession. That is, a shot should only be taken if its expected point value exceeds the foreword-looking value of continuing the arrival process of shot opportunities.

The key identification challenge in most stopping problems is estimating properties of the data generating process using a sub-sample generated by endogenous choices. In our context, we have to estimate a lineup's shot *opportunity* distribution based on the shots we *actually observe*. Our identification strategy exploits the fact that as the shot clock winds down, a strategic team will increase their shooting frequencies, allowing us to quantify the resulting trade-off in efficiency and shooting frequency. That is, we us time remaining on the shot clock as an instrument. This instrument will only be relevant if teams understand the logic of the very problem we are studying to at least some degree and we show that all the lineups in our study easily pass this requirement. Our main estimation relies on a parametric model of the opportunity distribution and while our model passes over-identification tests, we are careful to discuss how model miss-specification could impact the interpretation of our results.

We focus our estimation on the most commonly played lineup for each of the thirty teams to ensure adequate sample sizes.[2] Overall, these lineups adopt reservation thresholds, throughout the entire range of the shot clock, close to the possession's continuation value. As compared to an optimal "fixed" threshold, where the team always shoots in the final period and uses a single, point maximizing threshold in all other periods, the lineups we study capture 84% of the gains of dynamic decision making. Only a single lineup captures less than 66% of these gains.

The deviations from optimality we observe, while small overall, are statistically significant for 12 of 30 lineups (though only 4 after applying a Bonferroni adjustment for multiple hypothesis testing). We categorize mistakes along two dimensions. The first is a level shift of the reservation threshold, indicating a fixed propensity to be too eager or too reluctant to shoot. The second is an incorrect slope, indicating under- or over-responsiveness to time remaining. Both classes of mistakes are present in our data, but they go in a consistent direction. When the level is off, the threshold is too low. When the slope is off, it is too steep. Both push a lineup towards taking too many shots in the intermediate periods of the shot clock ("impatience"). The lineups with the least playing experience (with each other) tend to be the most impatient. Since coaches allocate playing timing based on a lineup's performance, we cannot determine whether this correlation represents learning or selection.[3]

In our model, shot selection is a group decision problem, but in reality only one player holds the ball at any given time. As such, the departures from optimal play could be driven by a failure of incentives. For example, a player could knowingly use a lower-than-optimal threshold in order increase individual point output per-minute while lowering his shooting percentage (accuracy). This strategy would increase "points-per-game" provided poor shot selection did not lead to reduced playing time. If teams focus on raw points-per-game when determining player salaries, then a failure to enforce group incentives would push towards impatient shooting.[4] Alternatively, the biases towards early stopping could be evidence of cognitive limitations. Early stopping has been repeatedly observed in laboratory studies (Seale and Rapoport, 2000; Zwick et al., 2003; Bearden et al., 2005, 2006), but lab subjects tend to exhibit excess flatness, not steepness, in their reservation thresholds. The fact that lineups with greater collective playing experience more closely adhere to optimal play is not dispositive either. In favor of the incentives story, this could indicate that coaches punish sub-optimal play. In favor of psychological causes, it could be viewed as evidence of learning. Ultimately we cannot distinguish the underlying causes of mistakes, but it is important to keep in the mind that the magnitude these errors are relatively small.

There is a deep experimental literature investigating subjects' ability to effectively use continuation values in search problems. The earliest papers studied the "classical secretary problem" in which a subject observes a series of "applicants" and wins a prize if she "hires" the best one. A robust finding is that subjects "under-search" by hiring a candidate too quickly (Kahan et al., 1967; Rapoport and Tversky, 1970). Recent work has relaxed many assumptions of the original formulation and has also found early stopping (Seale and Rapoport, 2000; Zwick et al., 2003; Bearden et al., 2005, 2006). Most relevant, when payments are simply the chosen draw, the majority of subjects incorrectly use flat reservation thresholds (Lee et al., 2004), even when the time horizon is short (Lee, 2006).[5] A similar pattern of implicitly undervaluing continuation values arises in alternating offer bargaining games(Ochs and Roth, 1989), a pattern that has been linked to a "limited lookahead" strategy (Johnson et al., 2002). In repeated prisoner's dilemmas, subjects are also less responsive to continuation probability than

[1] "Using a scoring opportunity" can mean taking a shot or it can mean making an aggressive play that may lead to either a turnover or a shot.

[2] While this does induce some selection as it leaves us with a team's better players and lineups with higher than average experience playing together, it has the benefit of focusing our analysis on the most economically important players and mitigating learning considerations.

[3] The correlation is present conditional on the other observable features of the lineup such as player quality or heterogeneity in player quality.

[4] Conversely, if too much weight is placed on shooting percentage, players would adopt too high or too flat a threshold. While NBA teams employ statisticians to evaluate player quality with advanced metrics, relatively small, selfishly motivated departures from optimal play could go undetected. It would thus be a stretch to assume an efficient market.

[5] With a very long time horizon, a high constant threshold does quite well since one will get a good draw with high likelihood. The short horizon games are a better analog to basketball.