


Stop the Clock: Are Timeout Effects Real?

Niander Assis , Renato Assunção, and Pedro O.S. Vaz-de-Melo

Departamento de Ciência da Computação
Universidade Federal de Minas Gerais, Belo Horizonte, Brazil
{niander,assuncao,olmo}@dcc.ufmg.br

Abstract. Timeout is a short interruption during games used to communicate a change in strategy, to give the players a rest or to stop a negative flow in the game. Whatever the reason, coaches expect an improvement in their team’s performance after a timeout. But how effective are these timeouts in doing so? The simple average of the differences between the scores before and after the timeouts has been used as evidence that there is an effect and that it is substantial. We claim that these statistical averages are not proper evidence and a more sound approach is needed. We applied a formal causal framework using a large dataset of official NBA play-by-play tables and drew our assumptions about the data generation process in a causal graph. Using different matching techniques to estimate the causal effect of timeouts, we concluded that timeouts have no effect on teams’ performances. Actually, since most timeouts are called when the opposing team is scoring more frequently, the moments that follow resemble an improvement in the team’s performance but are just the natural game tendency to return to its average state. This is another example of what statisticians call the *regression to the mean* phenomenon.

Keywords: causal inference · sports analytics · timeout effect · momentum · bayesian networks

1 Introduction

In sports, timeout is a short interruption in a play commonly used to stop a negative flow in the game, to discuss a strategy change, or to rest the players [3]. As this is the most direct way coaches can intervene during a game, their influence and strategic ability is best expressed during these events. A timeout is usually called when a team has a rather long streak of score losses [7, 24]. Popular belief and research [5, 11, 16, 18, 21] have found a positive effect on teams’ performances after the timeout. That is, on average, the team asking for the timeout recovers from the losses by scoring positively immediately after. This observed difference has been wrongly used as evidence that the timeout has a real and positive effect on teams’ performance. In order to answer such causal question, a formal counterfactual analysis should be used, and that is what we propose in this work.

There is an intense interest on causal models to analyze non-experimental data since causal reasoning can answer questions that machine learning itself

cannot [15]. Our approach is built on top of these causal inference approaches that are briefly introduced in Section 3. For each timeout event at time t_r in the database, we found a paired moment t_c in the same game when no timeout has been called that serves as a *control* moment for t_r , reflecting what would have happened the timeout had not been called. This control moment is chosen based on other variables about the current game instant, which were drawn in a causal graph that depict our assumptions about the generation of the data and, as we will further discuss, assess if the causal effect can be estimated without bias. In order to quantify how the game changed just after a given moment t , we proposed *Short-term Momentum Change* (STMC), which is discussed in Section 4.1.

After using a *matching* approach to construct our matched data (pairs of (t_r, t_c)), we found virtually no difference between the distribution of the STMC for real timeouts and control instants, i.e., the estimated *timeout effect* is very close to zero or non-existent. Hence, we conclude that the apparent positive effect of timeouts is another example of the well-known *regression to the mean fallacy* [1]. The dynamic match score fluctuates naturally and, after an intense increase, commonly returns towards a mild variation. Thus, because timeouts are usually called near the extreme moments, as we will show, the game seems to benefit to those losing. In summary, the main contributions of this paper are the following:

- We proposed a metric called *Short-term Momentum Change* (STMC) to quantify how much the game momentum changes after a time moment t_r associated with an event, such as a regular ball possession or an interruption of the game;
- We collected and organized a large dataset covering all the play-by-play information for all National Basketball Association (NBA) games of the four regular seasons from 2015 to 2018. A single season has over 280 thousands *game instants*, as we define in Section 4, and over 17 thousand timeout events;
- After a detailed causal inference analysis to evaluate the timeout effect of the *Short-term Momentum Change*, we did not find evidence that the timeout effect exists or that its effect size is meaningful. Inspired by others, we did also consider two other settings in which the timeout effect could be different: (i) only the last five minutes of the games and (ii) everything but the last five minutes.

The next section describes the previous work carried out on the effect of timeouts and discusses how our work distance from them. Section 3 gives a background on causal inference and the statistical models adopted. We start Section 4 by summarizing timeout rules in the NBA and describing our dataset, the play-by-play tables. In Section 4.1 we introduce our outcome variable of interest, the *Short-term Momentum Change*, and in Section 4.2 our causal model. In Section 4.3, we describe our treatment and control groups and, in Section 4.4, we explain our matching approaches. All the results are presented in Section 5. We close the paper in Section 6 with our conclusions.

2 Related Work

Timeouts are used and implemented in team sports for several reasons, such as to rest or change players, to inspire morale, to discuss plays, or to change the game strategy [20]. However, timeouts are mostly used to stop a negative flow in the game [7, 24], which is popularly referred as “the game momentum.” In basketball, momentum arises when one team is scoring significantly more than the other [11, 22].

Several earlier studies analyzed the effect timeouts have for decreasing the opponent’s momentum in the game [5, 11, 18, 21]. These studies analyze the effect of timeouts on teams’ performance just before and after it was called. For instance, by using a small sample of seven televised games from the 1989 National Collegiate Athletic Association (NCAA) tournament, Mace *et al.* [11] recorded specific events of interest, which were classified as either *Reinforcers* (e.g. successful shots) or *Adversities* (e.g. turnovers) and verified that the rate of these events change significantly among teams in the 3 minutes before and after each timeout. They found that while the team that called the timeout improved its performance, the opponent team decreased it. Other works reached the same conclusions using similar methodologies and different data sets [5, 18, 21].

To the best of our knowledge, Permutt [16] was the first to acknowledge the *regression to the mean* phenomenon in such analysis. Permutt considered specific game moments—timeouts called for after a team suffered a loss of six consecutive points. Similar to others, the short-term scoring ratio was observed to be higher after timeouts. However, in contrast to others, the paper compares real timeouts with other similar game moments without a timeout. With such analysis, Permutt found that timeouts can be effective at enhancing performance, but at a small magnitude. The most significant result shows that the home-team with a “first-half restriction” presents a 0.21 increase in average ratio for the next ten points. Calling a timeout predicts that the home-team will score 5.47 out of the next ten points as opposed to 5.26 points when a timeout is not called. Thus, the conclusion is that timeouts do not have any significant effect in changing the momentum of a game, i.e., using 6-0 runs as an indicator of instances where momentum would be a factor, teams were successful at “reversing” momentum even without the timeout as a mediator.

Although the work of Permutt [16] innovates by considering counterfactuals, the analyses still leave room for reasonable doubts about the reality of the timeout effect. It fails to take into account the existence of other important factors that could also influence on the momentum change and confound the true timeout effect. As a result, spurious correlations could have caused the lack of effect observed in the data. In our work, we take into account other factors such as coaches’ and team’s abilities, stadium and match conditions, clock time, quarter and relative score between the teams. More important, different from all the studies described in this section, we adopt a formally defined causal model approach [14] with a counterfactual analysis with its constructed control group. We show in a compelling way that timeouts do not have an effect in teams’ performances.

3 The Causality Framework

For illustration, consider Y our outcome variable and $A \in \{0, 1\}$ the *treatment variable*. Regardless of the actual value of A , we define $Y_{A=1}$ to be the value of Y had A been set to $A = 1$ and $Y_{A=0}$ to be the value of Y had A been set to $A = 0$. We say there is a causal effect of A on Y if $Y_{A=0} \neq Y_{A=1}$ and, conversely, there is no causal effect or the effect is null if $Y_{A=0} = Y_{A=1}$. These defined values are called *potential outcomes* [12] because just one potential outcome is factual, truly observed, while the others are counterfactuals. Therefore, we cannot generally identify the causal effects of a single individual. This problem is known as the *Fundamental Problem of Causal Inference (FPCI)* [8].

Nevertheless, in most causal inference settings, the real interest is in the population level effect, or the average causal effect defined by $\mathbb{E}(Y_{A=1} - Y_{A=0})$. $\mathbb{E}(Y | A = 1) - \mathbb{E}(Y | A = 0)$ gives a reliable estimate in randomized experimental studies, where treatment A is assigned randomly to each unit. However, in observational ones, we need to collect more information to control for and to make assumptions. One important assumption is the *conditional ignorability* [19]. This assumption is satisfied if, given a vector of covariates \mathbf{X} , the treatment variable A is conditionally independent of the potential outcomes ($Y_{A=0} \perp\!\!\!\perp A | \mathbf{X}$ and $Y_{A=1} \perp\!\!\!\perp A | \mathbf{X}$), and there is a positive probability of receiving treatment for all values of \mathbf{X} ($0 < \mathbb{P}(A = 1 | \mathbf{X} = \mathbf{x}) < 1$ for all \mathbf{x}). The conditional ignorability assumption allows us to state that $\mathbb{E}(Y_{A=1} - Y_{A=0} | \mathbf{X}) = \mathbb{E}(Y | A = 1, \mathbf{X}) - \mathbb{E}(Y | A = 0, \mathbf{X})$ for every value of \mathbf{X} . This represents the basic rationale behind the *matching* technique.

The simplest matching technique is the *exact matching*. For each possible $\mathbf{X} = \mathbf{x}$, we form two subgroups: one composed by individuals that received the treatment and have $\mathbf{X} = \mathbf{x}$, and the other by individuals that did not receive the treatment and have $\mathbf{X} = \mathbf{x}$. Unfortunately, exact matching is not feasible when the number of covariates is large or some are continuous. As an alternative, examples are usually matched according to a distance metric $d_{ij} = d(\mathbf{x}_i, \mathbf{x}_j)$ between the covariate configurations of pairs (i, j) of observations. The *Mahalanobis distance* [23] is a common choice for such distance metric as it takes into account the correlation between the different features in the vector \mathbf{X} . Another option is to use *propensity score* [19] for estimation of causal effects, which is defined as probability of receiving treatment given the covariates, i.e., $s(\mathbf{X}) = \mathbb{P}(A = 1 | \mathbf{X})$. Rosenbaum and Rubin proved that it is enough to just match on a distance calculated using the scalar scores $s(\mathbf{x})$, rather than the entire vector \mathbf{x} [19].

In general, each matching approach can be implemented using algorithms that are mainly classified as either *greedy* or *optimal*. The *greedy* ones, also known as *nearest-neighbor matching*, matches the i -th treated example with the available control example j that has the smallest distance d_{ij} . *optimal matching*, however, takes into account the whole reservoir of examples since the goal is to generate a matched sample that minimizes the total sum of distances between the pairs. Such optimal approach can be preferred in situations where there are great competition for controls. For a good review on different matching methods for causal inference, see [23].

Whatever matching technique used, its success can be partially judged by how balanced out are the covariates in the treatment and control groups. By pruning unmatched examples from the dataset, the control and treated groups of the remaining matched sample should have similar covariate distributions, when we say that matching achieves *balance* of the covariates distribution.

4 The Causal Effect of Timeout

According to the NBA 2016-2017 season official rules, in a professional NBA regular game, each team is entitled to six full-length timeouts and one 20-second timeout for each half. A full-length timeout can be of 60 seconds or 100 seconds, depending on when the timeout was requested. Also, every game has four regular periods plus the amount of overtime periods necessary on the occurrence of ties. There is a specific amount of timeouts expected in each period for commercial purposes. If neither team calls a timeout before a specific time, thus not fulfilling the next expected timeout, the official scorer stops the game and calls a timeout. The timeout is charged to the team that has not been charged before, starting with the *home* team. These timeouts are called *mandatory* or *official* timeouts.

In basketball games, possessions are new opportunities to score in the game. Each possession starts from the moment a team gets hold of the ball until one of his players scores, commits a fault, or loses the ball in defensive rebounds or turnovers. The total number of possessions are guaranteed to be approximately the same for both teams at the end of a match, so it provides a good standardization for the points scored by each team [9]. Indeed, most of basketball statistics are already given in a per possession manner.

Play-by-play tables capture the main play events such as goal attempts, rebounds, turnovers, faults, substitutions, timeouts and end of quarters (periods). For each play, we have the time in which it happened, the players and/or the team involved and any other relevant information, e.g., the score just after the play is recorded. Each play event is recorded as a new line in the table. While ball possessions are not clearly recorded in play-by-play tables, one can identify every change of possession from observing the game events.

In this work, we use play-by-play tables to identify the ball possessions and use them to observe how the teams' performance change when timeouts are called. We have identified every change of possession alongside the main interruptions—timeouts and end of quarters—in each game. Every change of possession and every main interruption is considered a new *game instant*. We model each basketball game as a series of discrete *game instants*. More formally, a *game instant* is either (1) a regular ball possession; or (2) a major game interruption, which can be a regular timeout, an official scorer timeout, or the end of a quarter. Player substitutions and fouls were not considered as a main interruption. In fact, substitutions can happen when the ball stops and not only during timeouts.

4.1 Short-term Momentum Change

Here we describe our outcome variable associated with teams' performance, for which we aim to estimate how it is affected by the timeouts. Let $\{P_t\}$ be an univariate stochastic process associated to a single match and indexed by the discrete game instants t . At any timeout moment, the team calling the timeout is defined as the *target team*. The P_t random variable at the end of the t -th game instant is the score of the target team minus the opposing team's score and it is called the scoring margin. Hence, P_t is a positive quantity when the target team is winning the match at game instant t and negative otherwise. At the end of the t -th game instant, $P_t = P_{t-1}$ in two situations. First, if t is a regular ball possession instant whose attacking team did not score, i.e., the possession ended with a turnover or defensive rebound. Second, if t is a main interruption instant and, consequently, none of the teams had the opportunity to score.

In order to evaluate how "momentum" changes after a game instant, we use the *Short-term Momentum Change (STMC)*, which is the amount by which the *scoring margin per possession rate* changes right after an game instant. For any game instant t and a positive integer $\lambda > 0$, we define the STMC, Y_t^λ , as the the average rate of change from P_t to $P_{t+\lambda}$ ($\Delta P_t^{t+\lambda}$) minus the average rate of change from $P_{t-\lambda-1}$ to P_{t-1} ($\Delta P_{t-\lambda}^t$). Note that we do not take into account the possible change in scoring margin caused in game instant t (the change from P_{t-1} to P_t):

$$Y_t^\lambda = \frac{P_{t+\lambda} - P_t}{\lambda} - \frac{P_{t-1} - P_{t-\lambda-1}}{\lambda} = \Delta P_t^{t+\lambda} - \Delta P_{t-\lambda}^t \quad (1)$$

for $t - \lambda \geq 0$ and $t + \lambda \leq n$, where n is the total number of game instants in a given game.

The hyper-parameter λ controls the time window used to evaluated how the game scoring dynamics changes around t . To balance out the offensive and defensive ball possessions, λ must be an even integer. Also, the variable Y_t^λ should only be evaluated if the interval $[t - \lambda, t + \lambda]$ contains no game interruptions, with the possible exception of t . In a causal perspective, λ represents our assumption for how many game instants that the interference (calling a timeout or not) at game instant t can influence and is influenced by, in the short-term.

Let A_t be the binary indicator that a timeout has been called at time t . We will denote $A_t = 1$ if a team calls a timeout right before the game instant t and $A_t = 0$ if t is a regular ball possession. If we find the set of covariates \mathbf{X} that satisfy the conditional ignorability assumption, we can apply a matching technique and our average causal effect of interest, $\mathbb{E}(Y_{A_t=1}^\lambda - Y_{A_t=0}^\lambda)$, can be estimated taking the difference in means from the matched treatment and control groups. The estimated *timeout effect TE* is defined as:

$$TE = \mathbb{E}(Y_t^\lambda | A_t = 1) - \mathbb{E}(Y_t^\lambda | A_t = 0). \quad (2)$$

Every game is composed by two teams, the *home* and the *away* team. Because we want to estimate the causal effect of timeouts on the performance of the *team that actually asked for it*, we decided to estimate the average causal effect of




Fig. 1. A causal graph to model the timeout effect.

timeouts called by the *home* teams (TE_h) and the *away* teams (TE_a), separately. We proceed now to present our causal model which encodes our assumptions.

4.2 The Causal Model

Pearl [14] suggests the use of directed acyclic graphs (DAGs) as a way of encoding causal model assumptions with nodes representing the random variables and the direct edges representing direct causal relationships. One can identify in such graph a set of variables (or nodes) that satisfies the so called *back-door criterion* [14]. These are variables that blocks all back-door paths from A (the treatment variable) to Y (the outcome variable) and does not include any descendants of A . Given that the graphical model includes all important confounding variables, it can be shown that conditioning on them suffices to remove all non-causal dependencies between A and Y . In other words, it leaves only causal dependence that corresponds to the causal effect.

There are many factors that can potentially influence the short-term performance (STMC) of the team that called a timeout after a given game instant. These can be *intra-game* factors, which vary along the game, such as the scoring margin, the quarter and the time since the start of the quarter, or *inter-game* factors, which vary from game to game, such as the venue conditions, the attendance at the venue, the specific adversary team, the players available and the teams' momentum in the season.

It is very intuitive why *intra-game* factors, which are specific to a game instant, are considered a cause of both the treatment and outcome, thus being considered a confounder. In a not so straightforward way, some *inter-game* factors are also very likely to affect both the treatment and outcome. For example, a team playing against a stronger or a weaker adversary would differently request the available timeouts and the afterwards performance may be differently affected.

Figure 1 shows our causal model graph. Each game instant t can either receive the treatment assignment $A_t = 1$ or $A_t = 0$, meaning that the game instant is a timeout or a regular ball possession, respectively. The variables \mathbf{X}_t represent the

observed covariates, which are *intra-game* factors specific to the game instant t : (i) the current quarter (period) (Q_t), (ii) the current scoring margin (P_t) and (iii) the current time in seconds since the start of the period (S_t). The variables represented by the node \mathbf{U} are the *inter-game* factors, or the covariates related to a *specific game* that influence both the treatment assignment and the game outcome as exemplified in the last paragraph. Most of these variables are not directly observed or very difficult to measure—players and coach strategies, teams’ relative skill difference and venue conditions. Hence, we include them in our graph as a dashed circle. The average rates of scoring margin change before ($\Delta P_{t-\lambda}^t$) and after ($\Delta P_t^{t+\lambda}$) the game instant t are also in the graph, as well as the outcome Y_t^λ that is connected by dashed edges since it is a deterministic node—a logical function of the other two stochastic nodes.

We are interested in the causal effect of A_t on Y_t^λ . Since $\Delta P_{t-\lambda}^t$ is a direct cause of A_t and not the reverse—for obvious chronological reasons—we actually want to measure the causal effect of A_t on $\Delta P_t^{t+\lambda}$. According to the *back-door* criterion [14], if we adjust for \mathbf{U} , \mathbf{X}_t and $\Delta P_{t-\lambda}^t$ we block any non-causal influence of A_t on $\Delta P_t^{t+\lambda}$.

4.3 Data

Because we want to estimate TE as defined in Equation (2), our treated and control groups are formed by game instants’ STMC, Y_t^λ . As discussed in Section 4.1, depending on which value we choose for λ , Y_t^λ is not valid—if the short-term window induced by λ includes another major interruption besides the possible t or is longer than the start or end of the game. Therefore, our inclusion criteria for both groups is that STMC exists and can be calculated. For a given λ , the treated group, $\{Y_t^\lambda | A_t = 1\}$, is formed by the valid real timeouts’ STMC. The control group, $\{Y_t^\lambda | A_t = 0\}$, is formed by any valid game instant t ’s STMC that is not a timeout or any other kind of major interruption.

Since we want to estimate TE_h and TE_a , we have two treatment groups, one for timeouts called by *home* teams and one for timeouts called by *away* teams. On the other hand, it does not make sense to classify the control group as either *home* or *away*, thus we have just one control group. We will limit ourselves in the future to just mention these treatment groups as either the *home* treatment group or *away* treatment group.

Our data consist of play-by-play information for every game from the 2014-2015, 2015-2016, 2016-2017 and 2017-2018 National Basketball Association NBA regular seasons. We crawled the data from the Basketball-Reference website¹. Most of our analysis will consist only of games from the NBA 2016-2017 season because using more than a single season would lead to very big samples that are impractical to apply our matching approaches. Also, while we did perform the same analysis using only other seasons, achieving very similar results, the choice for the 2016-2017 season is arbitrary, mainly due to be the first season for which we collected the data.

¹ <http://www.basketball-reference.com>

The 2016-2017 season had a total of 30 teams and 1,309 games (1,230 for the regular season and 79 in the playoffs). Considering all games, we computed 281,373 game instants, including the 17,765 identified timeouts (7,754 were called by *home* teams and 8,011 by *away* teams), and the 2,000 *mandatory* timeouts. Our datasets, code and further instructions on how to reproduce our results can be found at our GitHub repository².

4.4 Matching

The variables \mathbf{U} , \mathbf{X}_t and $\Delta P_{t-\lambda}^t$ should be controlled for. In other words, they should be considered as possible confounders. Consequently, all of these variables are included in our matching for a valid causal inference. While we consider \mathbf{U} , the *inter-game factors*, unobserved covariates, we can still control them by pairing timeout examples with non-timeout examples *taken from the same game*. Furthermore, the variable $\Delta P_{t-\lambda}^t$ is likely the most important confounder covariate in our model. Indeed, coaches tend to call a timeout when their teams are suffering from a bad “momentum”, evidencing great influence on the treatment assignment A_t . Also, $\Delta P_t^{t+\lambda}$, the average rate of scoring margin change after a game instant t , should be highly causal dependent on $\Delta P_{t-\lambda}^t$. Therefore, in whatever matching approach we use, timeouts and control examples *taken from the same game* and with equal $\Delta P_{t-\lambda}^t$ are going to be matched, hopefully, achieving balance for \mathbf{X}_t . We also restrict our matches to be constructed with *non-overlapping ball possessions*. This restriction arises from our assumption that λ defines a range of game instants that are dependent and influence A_t as discussed in Section 4.1.

We applied three matching procedures: (1) *no-balance matching*; (2) *Mahalanobis matching*, and (3) *propensity score matching*. In the no-balance matching, each treatment example is paired with a valid control example that has the same $\Delta P_{t-\lambda}^t$ and is taken from the same game. We did not consider \mathbf{X}_t in this matching. For the Mahalanobis matching, we applied the Mahalanobis distance using all covariates in \mathbf{X}_t , i.e., current quarter (Q_t), current scoring margin (P_t), and current clock time in seconds (S_t) since the start of the quarter. Finally, for the propensity score matching technique, we applied a simple euclidean distance match on the estimated scalar propensity score.

The true propensity score $s(\mathbf{X}) = \mathbb{P}(A = 1 | \mathbf{X})$ is unknown and must be estimated. Since estimating $\mathbb{P}(A = 1 | \mathbf{X})$ can be seen as a classification task, any a supervised classification model could be used. While logistic regression is the most common estimation procedure for propensity score, Lee *et al.* [10] showed that, in a non-linear dependence scenario, the use of machine learning models such as boosting regression trees to estimate the propensity score achieves better covariate balance in the matched sample. Indeed, our treatment assignment present a great non-linear dependence on its covariates. Take the clock time S_t , for example. As explained in Section 4.3, the timeout rules of NBA stimulate coaches to call a timeout just before a mandatory timeout would have been called

² <https://github.com/pkdd-paper/paper667>

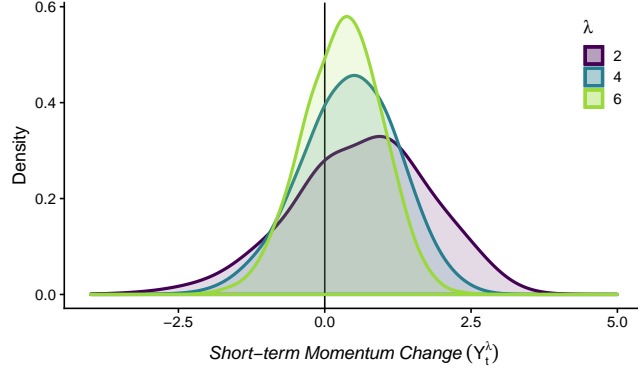


Fig. 2. The STMC (Y_t^λ) distribution for home and away timeouts, considering three different ball possession windows $\lambda = 2, 4, 6$.

by the official scorer. We use the boosting regression tree algorithm implemented in the *gbm* R package [6] to estimate the propensity score using \mathbf{X}_t .

Because we restrict timeout and non-timeout pairs to be taken from the same game, we have a very sparse matching problem. The *rcbalance* R package [17] implementation of *optimal matching* exploits such sparsity of treatment-control links to reduce computational time for larger problems. We use the optimal algorithm implemented in this package for all the aforementioned matching approaches. In addition, before applying any matching technique, we retained in the control subpopulation only those non-timeout game instants t' ($A_{t'} = 0$) for which the value $\Delta P_{t'-\lambda}^{t'}$ is exactly equal to at least one $\Delta P_{t-\lambda}^t$ calculated to a real timeout instant t ($A_t = 1$) in the same game. This improved the running performance even more.

5 Experimental Results

In order to find out whether our data shows the generally accepted positive correlation between timeouts and improvements in the “momentum”, we calculated Y_t^λ for every game instant associated with a timeout t in every single game using $\lambda = 2, 4, 6$. Figure 2 shows the estimated density distribution of the STMC Y_t^λ for all timeouts, including those called by both *home* and *away* teams, but removing the *official* timeouts. The sample means and number of valid timeout examples in each sample are 0.629 and 14,031 for $\lambda = 2$, 0.421 and 12,225 for $\lambda = 4$, 0.302 and 10,296 for $\lambda = 6$, respectively.

These results shows that, when a timeout is called by a team, its momentum improves by a small positive amount afterwards. For instance, with $\lambda = 4$, the average value of STMC is 0.421. That is, on average, there is an increase of 0.421 points for a team’s scoring margin per possession after it called the timeout. These results are consistent with previous works mentioned in Section 2 [11, 18, 16, 5, 21]. We applied the non-parametric ones-sample Wilcoxon statistical test

Table 1. Summary statistics and SMD for balance assessment for matching using *home* treatment group. The control ($A_t = 0$) and timeout ($A_t = 1$) groups are presented before (BM) and after all three matchings approaches: No-Balance (NB), Mahalanobis distance (M), and Propensity score (P)

λ	Method	S_t (mean(sd))			Q_t (mean(sd))			P_t (mean(sd))		
		$A_t = 0$	$A_t = 1$	SMD	$A_t = 0$	$A_t = 1$	SMD	$A_t = 0$	$A_t = 1$	SMD
2	BM	363.42 (198.77)	410.03 (168.47)	0.253	2.42 (1.12)	2.68 (1.15)	0.222	1.73 (10.81)	-0.00 (10.78)	0.161
	NB	363.11 (201.26)	410.20 (168.52)	0.254	2.47 (1.12)	2.68 (1.15)	0.178	0.56 (10.63)	0.00 (10.78)	0.052
	M	397.73 (168.83)	410.27 (168.50)	0.074	2.61 (1.11)	2.68 (1.15)	0.062	0.53 (10.86)	0.00 (10.78)	0.049
	P	403.85 (163.08)	410.14 (168.50)	0.038	2.67 (1.13)	2.68 (1.15)	0.009	0.23 (11.01)	0.00 (10.78)	0.021
4	BM	351.23 (187.75)	388.45 (153.17)	0.217	2.33 (1.11)	2.58 (1.13)	0.222	1.75 (10.69)	0.22 (11.04)	0.141
	NB	351.59 (191.15)	394.10 (151.23)	0.247	2.33 (1.11)	2.58 (1.14)	0.219	0.44 (10.44)	0.33 (11.14)	0.011
	M	381.35 (167.92)	393.79 (151.12)	0.078	2.43 (1.04)	2.57 (1.14)	0.135	0.50 (11.06)	0.32 (11.12)	0.016
	P	385.03 (162.81)	393.70 (151.37)	0.055	2.56 (1.11)	2.58 (1.14)	0.017	0.38 (11.33)	0.32 (11.12)	0.005
6	BM	334.96 (170.51)	380.80 (143.08)	0.291	2.19 (1.08)	2.49 (1.12)	0.270	1.67 (10.34)	0.33 (11.09)	0.124
	NB	332.33 (173.87)	389.79 (139.28)	0.365	2.20 (1.08)	2.48 (1.12)	0.253	0.71 (10.40)	0.79 (11.13)	0.007
	M	351.13 (162.90)	389.67 (138.95)	0.255	2.27 (1.02)	2.48 (1.12)	0.200	0.77 (10.80)	0.80 (11.17)	0.002
	P	356.99 (160.61)	389.72 (139.63)	0.218	2.35 (1.08)	2.48 (1.12)	0.122	0.72 (10.91)	0.81 (11.16)	0.008

and a bootstrap based test for the mean with the null hypothesis being that the mean is equal to zero. Both tests for the three different values of λ yielded p-values numerically equal to zero.

While these results could be used as evidence on why there is such common and widespread belief that timeouts improves teams' performance, or *breaks* the momentum, it is not an evidence of the causal effect of timeouts. We move on to consider the analysis under our causal framework discussed in Section 4.

5.1 Matching Results

Each of the three matching methods was applied twice: one time using the treatment group with *away* timeouts and the control group, and the other using the treatment group with *home* timeouts and the control group. Some examples did not find a valid match and, therefore, were not included in the matched samples. Also, it should be noted that all matches were performed without replacement.

To evaluate for proper covariate balance between the treatment groups, a common numerical discrepancy measurement is the difference in means divided by the pooled standard deviation of each covariate, known as the *standardized mean difference* (SMD) [4]. Unlike t-tests, SMD is not influenced by sample sizes and allows comparison between variables of different measured units. There is no general consensus on which value of SMD should denote an accepted imbalance level. Some researches, although, have proposed a threshold of 0.1 [13]. Table 1 summarizes the covariate distribution with its mean and standard deviation values and the SMD of our matched samples considering all different approaches.

For simplicity, we are only including here the results from the *home* timeouts matched samples. Indeed, the *away* samples showed very similar results and

it can also be accessible from our GitHub repository³. Also, we do not show balance for $\Delta P_{t-\lambda}^t$ as it is perfectly balanced due to our perfect match on this covariate. The unmatched sample sizes for control groups are 172, 785, 101, 093 and 49, 403; and, for treatment groups, 6, 912, 6, 048 and 5, 127 with $\lambda = 2, 4, 6$, respectively. For the matched samples, because of our 1:1 matching approach, the sample sizes are equal in both treatment and control groups, even across the different matching techniques—for $\lambda = 2, 4, 6$, they were 6, 895, 5, 477 and 3, 832, respectively.

From an initial look we can see that there are very similar results in terms of covariate balance for all three different matching approaches. Also, with no surprise, the matched treatment and control samples have equal sizes and the $\Delta P_{t-\lambda}^t$ covariate is equally distributed. Comparing the *no-balance* matching with the balance in *before matching*, it is clear that this simple matching procedure reduces substantially the SMD for all covariates. However, this is not enough to make the SMD negligible for all cases and hence whatever conclusions based on the comparison of these no-balance matched samples will be rightly subjected to doubt. Considering the *Mahalanobis matching*, it achieved better covariate balance for all covariates and λ values in comparison with *before matching* and *no-balance matching*. For $\lambda = 2$ and $\lambda = 4$, all SMDs are below 0.1, with the exception of the Q_t covariate in $\lambda = 4$, in both analysis. The $\lambda = 6$ configuration, on the other hand, presented the worse SMDs— P_t is the only covariate with SMD below the 0.1 mark. The *Mahalanobis matching* was not as good as the *Propensity score matching*. All covariates have SMDs smaller than 0.1 for $\lambda = 2$ and $\lambda = 4$ in both analysis. For $\lambda = 6$, while Q_t and S_t still had SMD above the 0.1 mark, it is still smaller than the ones obtained in the previous matching for the same configuration.

Plotting the covariate distributions from both treatment and control groups is a qualitative alternative of checking proper balance. Figure 3 shows the covariate distributions for S_t , Q_t and P_t after all three matchings with $\lambda = 4$ on the analysis using timeouts called by the *home* team. We can see that *No-balance* and *Mahalanobis* matching presented some imbalance for S_t and Q_t . On the other hand, *propensity score matching* shows rather similar distributions suggesting a much better balance. The fact that timeouts and control examples are matched only if taken from the same game makes it more complicated to find matches with a more balanced \mathbf{X}_t . *Mahalanobis matching*, while trying to match samples as close as possible, encounters great difficulties. Propensity score, however, translates all interactions and non-linearities presented in the joint distribution of all covariates with the propensity score. This is why its matching was better.

5.2 Timeout Effect

We analyzed our matched data using a Monte Carlo permutation test [2, Vol II, chapter 10] and difference of means between the two groups as our test statistic. When the null hypothesis is true, the timeout effect TE defined in (2) is equal to

³ <https://github.com/pkdd-paper/paper667>

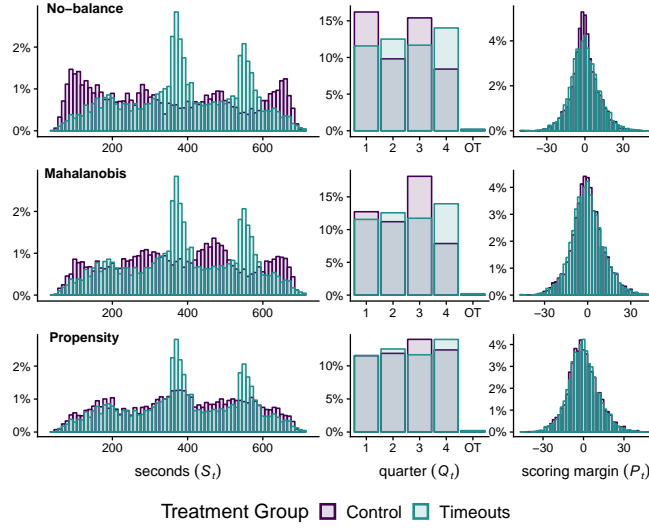


Fig. 3. The control (magenta) and timeout (turquoise) *seconds*, *quarter* and *scoring margin* covariate (\mathbf{X}_t) distributions after all matchings using the *home* timeouts treatment group and $\lambda = 4$.

zero. The treatment label was permuted ten thousand times. It should be noted that we have large samples—the number of treatment-control pairs in each test varies from 3,766 to 7,082. Hence, we believe our analysis would mostly benefit if effect sizes, alongside statistical significance, are taken into account for our conclusions. Table 2 shows the estimated average timeout effects for the *home* (TE_h) and the *away* (TE_a) teams. In addition, we have also included in the table a 99% level confidence interval generated from the p-values obtained with the Monte Carlo permutation test.

There are two fundamental remarks here. First, the estimated timeout effect TE is very small for all cases, practically irrelevant during a game. Remember that the timeout effect TE defined in Equation (2) is the amount by which the team should expect its scoring margin per possession to change if a timeout were to be called by them. It would be rather challenging to find the minimum effect size for which it would be deemed enough to change the teams performance. It would not be a bad idea, however, to consider it to be an effect of at least *one* marginal point. Specially because our assumption is that timeouts have an effect in the short-term window represented by λ . Yet, from our results, take the propensity matching with $\lambda = 4$, which yielded the largest absolute timeout effect equal to -0.059 , for instance. This means that if a team calls a timeout, it should expect its scoring margin to decrease 0.059 points per possession in the next 4 possessions, which can be considered negligible in a basketball game. Second, by analyzing the confidence intervals, the number of significant tests is very small given the very large number of examples in each of them. From the 18

Table 2. The estimated timeout effect for both *away* (TE_a) and *home* (TE_h) timeouts under different matching procedures and λ values. The respective confidence intervals for each timeout effect for 99% level is also shown. While some values are statistically significant, all timeout effects are negligible due to the small effect sizes.

λ Method	TE_a	99% CI	TE_h	99% CI
2 No-balance	-0.028	(-0.075, 0.020)	-0.022	(-0.072, 0.028)
Mahalanobis	-0.032	(-0.079, 0.015)	-0.017	(-0.067, 0.032)
Propensity	-0.044	(-0.092, 0.004)	-0.021	(-0.071, 0.028)
4 No-balance	-0.043	(-0.082, -0.004)	-0.023	(-0.063, 0.017)
Mahalanobis	-0.053	(-0.092, -0.013)	-0.013	(-0.053, 0.028)
Propensity	-0.059	(-0.098, -0.020)	-0.032	(-0.072, 0.008)
6 No-balance	-0.031	(-0.068, 0.005)	-0.036	(-0.072, 0.000)
Mahalanobis	-0.046	(-0.083, -0.010)	-0.036	(-0.072, 0.000)
Propensity	-0.044	(-0.081, -0.008)	-0.046	(-0.082, -0.009)

Table 3. Additional analysis for excluding or considering only the last 5 minutes. For each analysis, the same matching approaches were applied using the new subsets of data. The results are very similar to the original analysis considering the whole game.

λ Method	Minus Last 5 Min				Only Last 5 Min			
	TE_a	99% CI	TE_h	99% CI	TE_a	99% CI	TE_h	99% CI
2 No-balance	-0.029	(-0.082, 0.024)	-0.010	(-0.065, 0.045)	-0.083	(-0.115, -0.052)	-0.033	(-0.067, 0.001)
Mahalanobis	-0.021	(-0.074, 0.032)	-0.010	(-0.065, 0.045)	-0.091	(-0.122, -0.060)	-0.040	(-0.074, -0.006)
Propensity	-0.044	(-0.098, 0.009)	-0.012	(-0.067, 0.043)	-0.094	(-0.125, -0.063)	-0.038	(-0.072, -0.004)
4 No-balance	-0.036	(-0.078, 0.006)	-0.030	(-0.073, 0.013)	-0.006	(-0.068, 0.055)	0.038	(-0.029, 0.105)
Mahalanobis	-0.055	(-0.097, -0.013)	-0.022	(-0.065, 0.020)	-0.019	(-0.081, 0.044)	0.026	(-0.041, 0.093)
Propensity	-0.057	(-0.099, -0.015)	-0.038	(-0.081, 0.004)	-0.019	(-0.082, 0.043)	0.007	(-0.060, 0.074)
6 No-balance	-0.044	(-0.082, -0.005)	-0.041	(-0.079, -0.003)	-0.046	(-0.315, 0.223)	-0.032	(-0.250, 0.186)
Mahalanobis	-0.051	(-0.089, -0.012)	-0.040	(-0.078, -0.002)	-0.028	(-0.299, 0.244)	-0.048	(-0.267, 0.172)
Propensity	-0.046	(-0.085, -0.008)	-0.055	(-0.093, -0.017)	-0.046	(-0.317, 0.224)	-0.048	(-0.265, 0.170)

tests, only 6 were statistically significant in the $\alpha = 0.001$ level. In fact, these confidence intervals barely include the 0 effect value. Nevertheless, while these tests were statistically significant, they are still negligible effects. Also, we want to point it out that we did not perform any adjustment for multiplicity of tests. Indeed, such adjustment would yield a smaller number of statistically significant tests.

We went further and investigated the *timeout effect* TE for two particular cases: (i) when the last five minutes are excluded and (ii) when only the final five minutes of the games are taken into account. For both cases, we rerun the matching approaches on the new subsets of the data. Before executing the matching approaches, we filtered out from treatment and control groups examples that happened within the last five minutes of the last quarter (the 4th) of each game for (i). In a similar fashion, we filtered out from treatment and control groups examples that happened outside of the last five minutes of each game for (ii). However, because we ended up with fewer sample units available for

matching, we included examples extracted from the other NBA seasons for the (ii) case, i.e., the 2014-2015, 2015-2016 and 2017-2018 seasons.

The results from both of these new analysis are in Table 3. For the case of excluding the last five minutes, we can see that we have slightly more statistical significant tests, 8 out of 18. Still, all of them have small effect sizes, making them not practically significant. The largest estimated effect is -0.057 , which is found under the same configuration that we found the largest in the original analysis: TE_a for the propensity score matching with $\lambda = 4$. For the case of considering only the last five minutes, the only 5 statistically significant tests were all found with $\lambda = 2$, but again, with very small effect sizes. The largest absolute effect is -0.094 for propensity score matching.

6 Conclusion

In this work we proposed a causality framework to quantify the effect of timeouts on basketball games. For the best of our knowledge, we were the first to resort on the theory of causality to solve this problem. While all previous studies pointed to a positive timeout effect, by applying our causality model on a large dataset of official NBA play-by-play data, we concluded that timeouts have no effect on teams' performance. This is another example of what statisticians call the *regression to the mean* phenomenon. Since most timeouts are called when the opponent team is scoring more frequently, the moments that follow resemble an improvement in the team's performance, but are just the natural game tendency to return to its average state. We have also stratified our analysis by either including only the last five minutes or everything but the last five minutes of all games, but the results pointed to the same conclusion: timeouts have virtually no effect on team's performance.

Acknowledgments This work is supported by the authors' individual grants from FAPEMIG, CAPES and CNPq.

References

1. Barnett, A.G., Van Der Pols, J.C., Dobson, A.J.: Regression to the mean: what it is and how to deal with it. *International journal of epidemiology* **34**(1), 215–220 (2004)
2. Bickel, P.J., Doksum, K.A.: *Mathematical Statistics: Basic Ideas and Selected Topics*, Volumes I-II. Chapman and Hall/CRC (2015)
3. Coffino, M.J.: *Odds-On Basketball Coaching: Crafting High-Percentage Strategies for Game Situations*. Rowman & Littlefield Publishers (2017)
4. Flury, B.K., Riedwyl, H.: Standard distance in univariate and multivariate analysis. *The American Statistician* **40**(3), 249–251 (1986)
5. Gómez, M.A., Jiménez, S., Navarro, R., Lago-Penas, C., Sampaio, J.: Effects of coaches' timeouts on basketball teams' offensive and defensive performances according to momentary differences in score and game period. *European Journal of Sport Science* **11**(5), 303–308 (2011)

6. Greenwell, B., Boehmke, B., Cunningham, J., Developers, G.: *gbm: Generalized Boosted Regression Models* (2019), <https://CRAN.R-project.org/package=gbm>, r package version 2.1.5
7. Halldorsson, V.: Coaches Use of Team Timeouts in Handball: A Mixed Method Analysis. *The Open Sports Sciences Journal* **9**(1), 143–152 (oct 2016)
8. Holland, P.W.: Statistics and causal inference. *Journal of the American statistical Association* **81**(396), 945–960 (1986)
9. Kubatko, J., Oliver, D., Pelton, K., Rosenbaum, D.T.: A starting point for analyzing basketball statistics. *Journal of Quantitative Analysis in Sports* **3**(3) (2007)
10. Lee, B.K., Lessler, J., Stuart, E.A.: Improving propensity score weighting using machine learning. *Statistics in medicine* **29**(3), 337–346 (2010)
11. Mace, F.C., Lalli, J.S., Shea, M.C., Nevin, J.A.: Behavioral momentum in college basketball. *Journal of Applied Behavior Analysis* **25**(3), 657–663 (1992)
12. Neyman, J.: edited and translated by dorota m. dabrowska and terrence p. speed (1990). on the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science* **5**(4), 465–472 (1923)
13. Normand, S.L.T., Landrum, M.B., Guadagnoli, E., Ayanian, J.Z., Ryan, T.J., Cleary, P.D., McNeil, B.J.: Validating recommendations for coronary angiography following acute myocardial infarction in the elderly: a matched analysis using propensity scores. *Journal of clinical epidemiology* **54**(4), 387–398 (2001)
14. Pearl, J.: *Causality. Causality: Models, Reasoning, and Inference*, Cambridge University Press (2009)
15. Pearl, J., Mackenzie, D.: *The Book of why: The New Science of Cause and Effect*. Basic Books (2018)
16. Permutt, S.: *The Efficacy of Momentum-Stopping Timeouts on Short-Term Performance in the National Basketball Association*. Ph.D. thesis, Haverford College. Department of Economics (2011)
17. Pimentel, S.D.: *rebalance: Large, Sparse Optimal Matching with Refined Covariate Balance* (2017), <https://CRAN.R-project.org/package=rcbalance>, r package version 1.8.5
18. Roane, H.S., Kelley, M.E., Trosclair, N.M., Hauer, L.S.: Behavioral momentum in sports: a partial replication with women’s basketball. *Journal of Applied Behavior Analysis* **37**(3), 385–390 (2004)
19. Rosenbaum, P.R., Rubin, D.B.: The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**(1), 41–55 (1983)
20. Saavedra, S., Mukherjee, S., Bagrow, J.P.: Is coaching experience associated with effective use of timeouts in basketball? *Scientific reports* **2**, 676 (2012)
21. Sampaio, J., Lago-Peñas, C., Gómez, M.A.: Brief exploration of short and mid-term timeout effects on basketball scoring according to situational variables. *European Journal of Sport Science* **13**(1), 25–30 (2013)
22. Siva, J.M., Cornelius, A.E., Finch, L.M.: Psychological Momentum and Skill Performance: A Laboratory Study. *Journal of Sport and Exercise Psychology* **14**(2), 119–133 (mar 1992)
23. Stuart, E.A.: Matching methods for causal inference: A review and a look forward. *Statistical science* **25**(1), 1 (2010)
24. Zetou, E., Kourtesis, T., Giazitzi, K., Michalopoulou, M.: Management and Content Analysis of Timeout during Volleyball Games. *International Journal of Performance Analysis in Sport* **8**(1), 44–55 (feb 2008)