



Introduction

Athletes' performances improve, peak, and eventually decline. This curve is called the "age curve" and we expect this curve to have heterogeneity with respect to the characteristics of the players. In this work, we focus on estimating the effect of rest between games on performance for each age. This is helpful for making decisions about resting a player and so-called "load management". We make three main contributions: First, we construct a Conditional Expectation Function (CEF) to compare the age curve for different covariates and treatments. Second, using a causal inference approach, we propose a methodology to construct age conditioned treatment effect (ACTE) for a given treatment. The ACTE can test causal hypotheses for each age on the treatment and outcomes of interest. Third, we apply this method to assess the effect of days between games on multiple performance metrics conditional on age.

Notation, Framework, and Estimands

We assume a superpopulation or distribution \mathcal{P} where we draw N independent tuples $(Y_i, A_i, W_i, X_i) \sim \mathcal{P}$ where Y is the observed outcome of interest, $A \in \mathbb{Z}$ is the age, $W \in \{0, 1\}$ is the binary treatment of interest, and $X \in \mathcal{X}$ is the covariate matrix. Under Rubin's potential outcome framework [2] the Conditional Expectation Function (CEF) for a given age a is defined as

$$\mu_0(a, x) := \mathbb{E}[Y(0)|A = a, X = x] \text{ and } \mu_1(a, x) := \mathbb{E}[Y(1)|A = a, X = x]$$

The construction of the CEF helps us compare the age curve for different characteristics. The curve is helpful for decision-making for each franchise to plan on the contract of players such as position, role on the team, country of origin, or any other covariates of the players. This could help us remove the selection bias on the age curve since as noted in [3], the bias exists. The comparison of the age curve between different values of covariates,

Then the Age Conditional Treatment Effect (ACTE) is expressed as

$$ACTE(a, x) := \mathbb{E}[Y(1) - Y(0)|A = a, X = x] = \mu_1(a, x) - \mu_0(a, x)$$

The ACTE is the effect of the treatment W for a given age and other covariates of interest. This is helpful to plan the treatment allocation given the player's attributes.

To identify the ACTE, we need the following traditional causal inference assumptions:

- **Probabilistic Assignment:** The assignment probability is strictly between 0 and 1

$$0 < \mathbb{P}(W_i = 1|X_i, Y_i(0), Y_i(1)) < 1$$

- **Consistency:** There is only one version of the potential outcome for each w

$$Y_i^{obs} = Y_i(w_i)$$

- **Unconfoundedness:** There are no unmeasured factors affecting the probability of treatment allocation and the outcome.

$$(Y_i(0), Y_i(1)) \perp\!\!\!\perp W_i|X_i$$

To claim causality, we need to discuss whether these assumptions are sensible given the covariates we use to control.

S-learner/Splines

'S-learner is one of the causal ML meta-algorithms introduced by [1]. It is referred to as the S-learner because it uses a "single" estimator. In this learner, the treatment is included in the estimator and we denote the CEF as

$$\mu_0(a, x) := \mathbb{E}[Y(0)|A = a, X = x, W = 0] \text{ and } \mu_1(a, x) := \mathbb{E}[Y(1)|A = a, X = x, W = 1]$$

. To estimate μ , we could use any base learner on the entire dataset. In particular, we will use the age spline regression method as described in [3] as our base learner.

The advantage of the S-learner is that it is simple to fit. The disadvantage is that for more complicated structures of ACTE, other meta-learners approaches such as X-learner perform better.

X-learner/Splines

As our first step, estimate two learners for treatment and control

$$\mu_0(x, a) = \mathbb{E}[Y(0)|A = a, X = x], \text{ and } \mu_0(x, a) = \mathbb{E}[Y(0)|A = a, X = x]$$

using age spline regression. For the second step, we calculate the "bias" using the model we trained in step one and feed the data set that we did not use to train. Namely, we feed the control data into the model trained using the treated data and vice versa:

$$\tilde{D}^1 = Y^1 - \hat{\mu}_0(a, x^1) \text{ and } \tilde{D}^0 = Y^0 - \hat{\mu}_1(a, x^0)$$

. Next, we use age spline regressions to estimate $\tau^0(a, x) = \mathbb{E}[\tilde{D}^1|A = a, X = x]$ and $\tau^1(a, x) = \mathbb{E}[\tilde{D}^0|A = a, X = x]$ and call its estimates $\hat{\tau}^0(a, x)$ and $\hat{\tau}^1(a, x)$. Finally, the ACTE for each age and covariates is estimated by

$$\hat{\tau}(a, x) = g(a, x)\hat{\tau}_0(a, x) + (1 - g(a, x))\hat{\tau}_1(a, x)$$

where $g \in [0, 1]$ is a weight function. Typically, we use the propensity score.

The advantage of this method is that the curve constructed by this approach is smooth and easier to interpret than the popular X-learners using random forests or XGBoost. Additionally, thanks to the X-learner's property, it can handle more complex ACTE functions than the S-learner.

Application to NBA data: Assessing load-management in sports

Using 10 seasons (2011-2022) of NBA data, we apply the above methods to estimate the effect of days of rest and age on points per 100 possessions played for NBA players. The treatment $w = 0$ indicates back-to-back (b2b) games and $w = 1$ indicates 1+day(s) rest, $Y(w)$ is the potential outcome when treatment is w , and X would be the covariate matrix. Then $\mu_w(a, x)$ indicates the conditional expectation of the potential outcome $Y(w)$ given age a . This function allows us to construct the age curve. By using $\tau(a, x)$, we get the ACTE, which is the effect of resting over playing b2b games given age.

As this is an observational study, we might not be able to remove entirely the confounders but we will control the model using player, team, team against, home-away indicator, and season fixed effect.

Results: S-learner/Splines

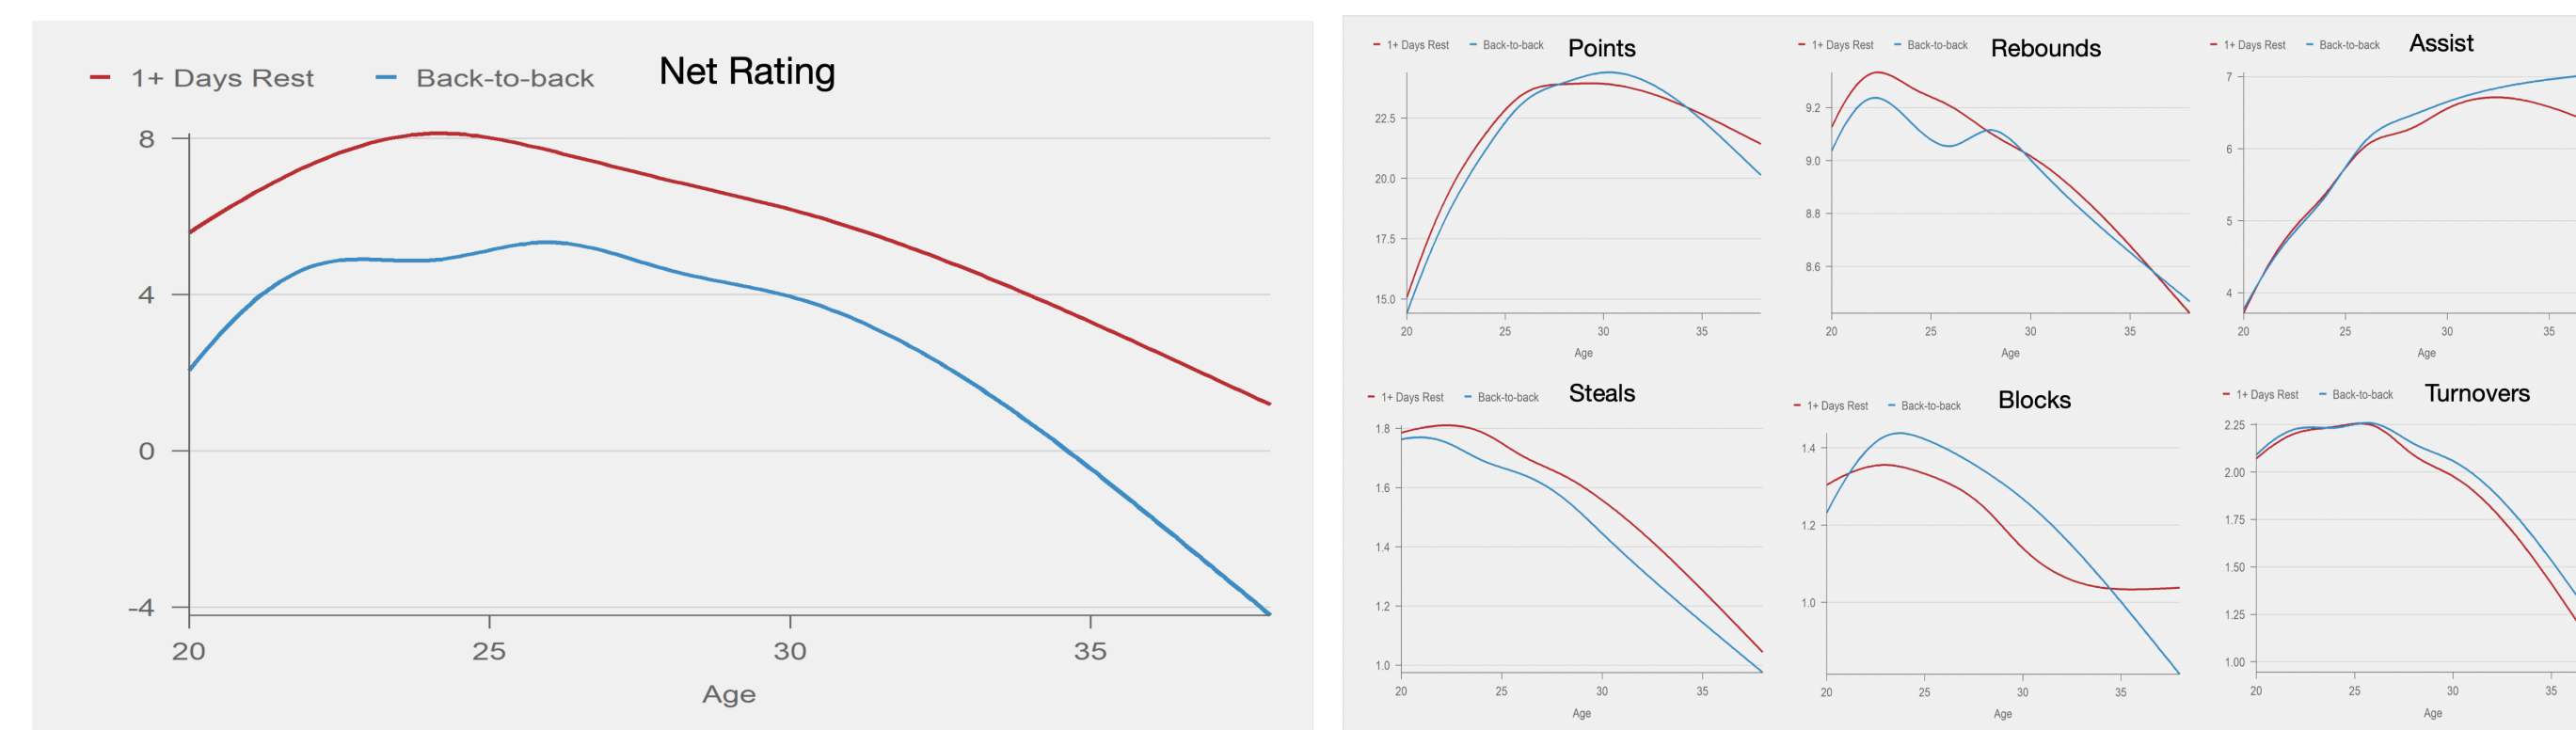


Figure 1. S-learner/Splines

The results of the S-learner can be observed in Figure 1. For net rating, we observe that b2b worsens the performance and moreover, the difference between the two curves increases as the player gets older.

We show the other metrics on the left. Points/100 poss. for ages considered as players' "peak" they successfully score without resting but for older ages, the scoring ability decreases faster. For rebounds, we observe that the scale is very small and we can conclude that there is a small difference between the two curves. Assists are interesting as players' do better for consecutive games. We believe that this happens as passing requires less physicality than scoring, an old player who played the game the day before might prefer passing rather than scoring due to fatigue. Resting is consistently better for steals and young players have higher steals per possession.

We believe the curve has a fitting problem for blocks as more than 60% of the players have 0 blocks. We could improve by fitting a Poisson regression instead of using the Gaussian link. Younger players make more turnovers and as they age, they tend to commit more turnovers in b2b games.

Results: X-learner/Splines

In Figure 2, we can observe the result of the X-learner using spline regressions.

On the left, we can observe the ACTE for each per 100 possessions variables. We can spot that points per 100 possessions are the most affected statistics using X-learner/Spline method. We are aware that the scaling for each of these variables might not be the same but for other outcomes except for points, the effect is marginal.

On the right, we observe that the net and offensive ratings have increasing trends. And the defensive rating is flat. We can observe that resting would affect offense more than defensive rating. Overall, the net rating improves and we notice that it is coming from the improvement in the offensive rating.

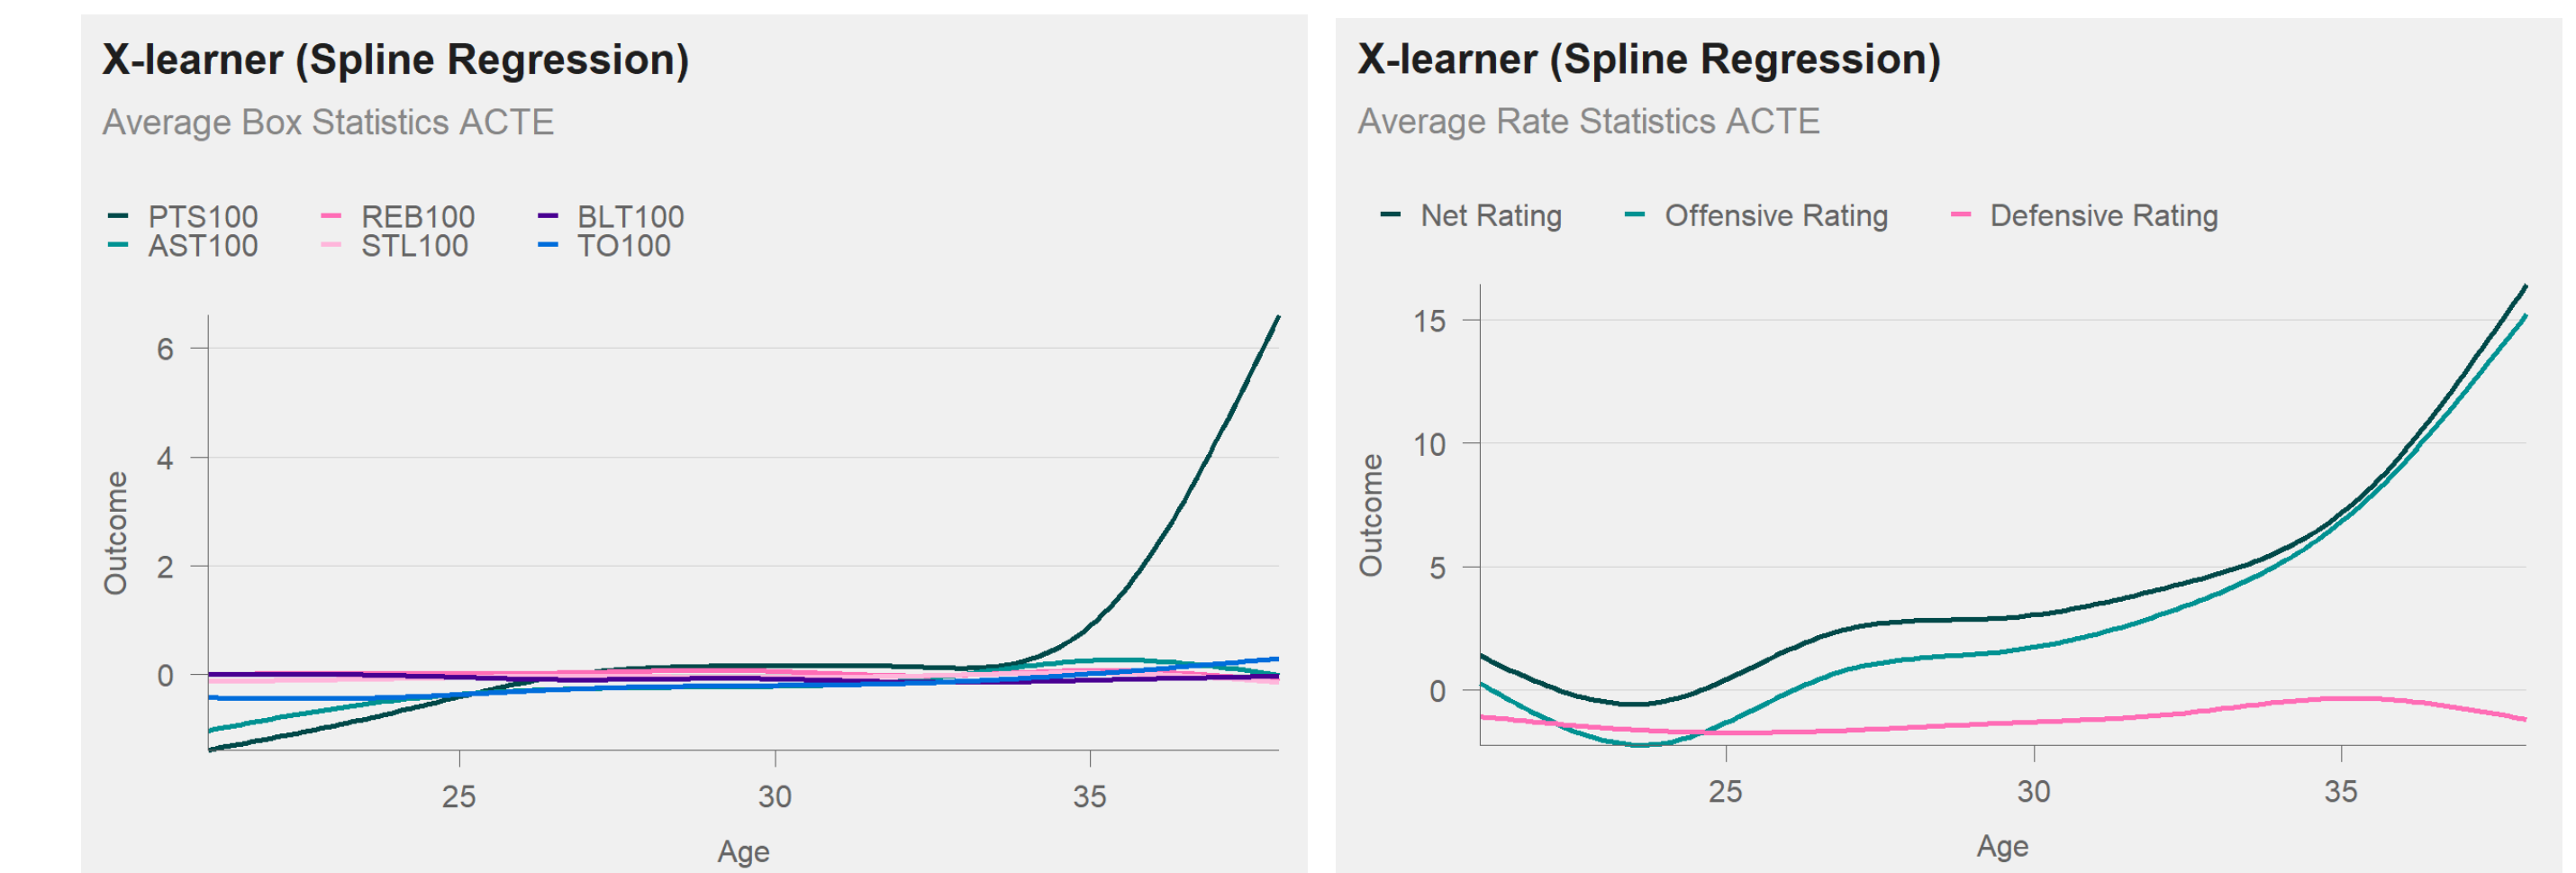


Figure 2. X-learners /Splines

S-learner vs X-learner

There is no right answer on which learner is better or worse. The advantage of S-learner is that we can plot the curve for both control and treatment and it is easy to interpret. On the other hand, S-learner has disadvantages, it can be biased towards zero, and when the ACTE is complex X-learner usually performs better. It might be safe to conclude that at least for points per 100 poss. and net rating there is a significant effect of resting as we observe this in both methods.

Discussion

In this work, we connected the age curve with the causal ML literature in this work. We applied this to NBA data to assess players' performance changes in b2b games. This framework can be applied in numerous settings to determine the hidden causality in the data. Although we used the model that best performed on [3] where they explored different techniques to account for this bias, we would like to solve from an instrumental variable perspective by trying to estimate the Survival Average Causal effect (SACE) using principal stratification.

References

- [1] Sören R. Künnel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, March 2019. doi:10.1073/pnas.1804597116. URL <https://www.pnas.org/doi/10.1073/pnas.1804597116>. Publisher: Proceedings of the National Academy of Sciences.
- [2] Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66:688–701, 1974. ISSN 1939-2176. doi:10.1037/h0037350. Place: US Publisher: American Psychological Association.
- [3] Michael Schuckers, Michael Lopez, and Brian Macdonald. What does not get observed can be used to make age curves stronger: estimating player age curves using regression and imputation, October 2021. URL <http://arxiv.org/abs/2110.14017>. arXiv:2110.14017 [stat].