If You Take The Road Less Travelled by, Does it Make a Difference?

Pittsburgh Penguins MSP Consulting Project

5/11/2021

Minyue Fan, Steve Kim, Kexiong Shen, Linda Yang

Advisor: Brian MacDonald

Client: Sam Ventura

Carnegie Mellon University

Agenda

Introduction: Overview

- There are multiple traditional paths hockey prospects can take to get to the NHL:
  - USHL → NCAA → NHL
  - USHL → NCAA → AHL → NHL
  - International → KHL → NHL
  - Other defined paths
- Most players do not immediately go to the NHL when they are eligible (drafted or not). They stay in or move to some “development leagues” before entering the NHL.
  - Draft eligibility (North America): Players must be 18 years old by 15 September and under 20 years old by 31 December in the year of the draft.
  - Development leagues: USHL, NCAA, etc.

Introduction: Research Question

- Questions:

  Does taking different development paths matter?

  How do players' development paths impact their performance and success in the NHL?

- The understanding in the scouting community is that development path does matter
  - Only anecdotal
  - We intend to establish grounding on this thought

4

Introduction: Overview

- People have very strong opinions about how players' development paths impact their future in the NHL.
    - Typically, American players who take the NCAA path have higher success rates (e.g. 20% make the NHL, compared to 5% from the USHL path)
    - However, the NCAA player pool are already better in terms of quality. Better players are getting their opportunities in the NCAA.
    - Is there causal impact of taking the NCAA path?

Carnegie Mellon University

Data

- Two datasets:
    - Leagues: NHL, NCAA, USHL and AHL
    - Time period: 2001 - 2020
        - *contains some data earlier than 2001*
    - Players’ biographical information
    - Players’ performance data each season
        - *box score statistics*

Data Description

- Biographical information:
  - 15786 players

	Player	Position	DateofBirth	Height	Weight	Nation	Shoots
1	Scott May	C	Jan 08, 1982	5'10" / 178 cm	187 lbs / 85 kg	Canada	R
2	Kent Gillings	F	Jun 14, 1979	5'10" / 177 cm	194 lbs / 88 kg	Canada / Ireland	R
3	Tyler Kindle	D	Feb 20, 1978	5'8" / 173 cm	165 lbs / 75 kg	USA	L
4	D'Arcy McConvey	C	Oct 23, 1981	5'10" / 177 cm	185 lbs / 84 kg	Canada	L
5	Lloyd Marks	C	Oct 21, 1977	5'8" / 173 cm	174 lbs / 79 kg	Canada	L
6	Jason Deskins	C	May 06, 1979	5'10" / 178 cm	185 lbs / 84 kg	USA	-
7	Jim Abbott	LW/C	May 03, 1980	6'1" / 186 cm	185 lbs / 84 kg	USA	L

- Rare missing data

Carnegie Mellon University

Data Description

- Player performance:
  - 266,326 rows (15,220 players * the number of seasons they played)

	Player	Season	Team	League	Games	Goals	Assists	TotalPoints	PenaltyMinutes	PlusMinus
1	Scott May	1998-99	South Surrey Eagles	BCHL	45	10	28	38	23	
2	Scott May	1999-00	South Surrey Eagles	BCHL	54	42	42	84	86	
3	Scott May	2000-01	Ohio State Univ.	NCAA	37	9	9	18	26	-3
4	Scott May	2001-02	Ohio State Univ.	NCAA	40	12	17	29	42	4
5	Scott May	2002-03	Ohio State Univ.	NCAA	43	10	25	35	56	5
6	Scott May	2003-04	Ohio State Univ.	NCAA	41	15	19	34	42	4
7	Scott May		St. John's Maple Leafs	AHL	5	1	1	2	2	3
8	Scott May	2004-05	St. John's Maple Leafs	AHL	16	0	1	1	21	-3

Exploratory Data Analysis

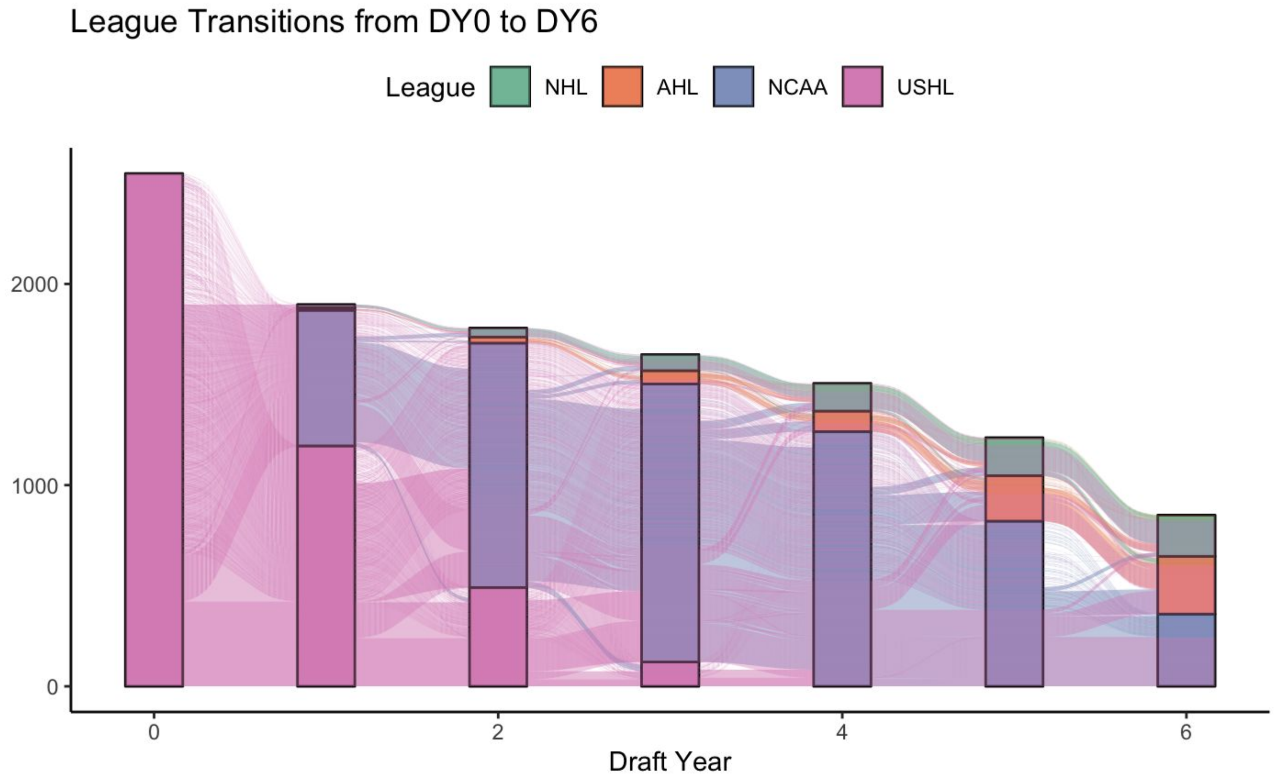

Carnegie Mellon University

Exploratory Data Analysis

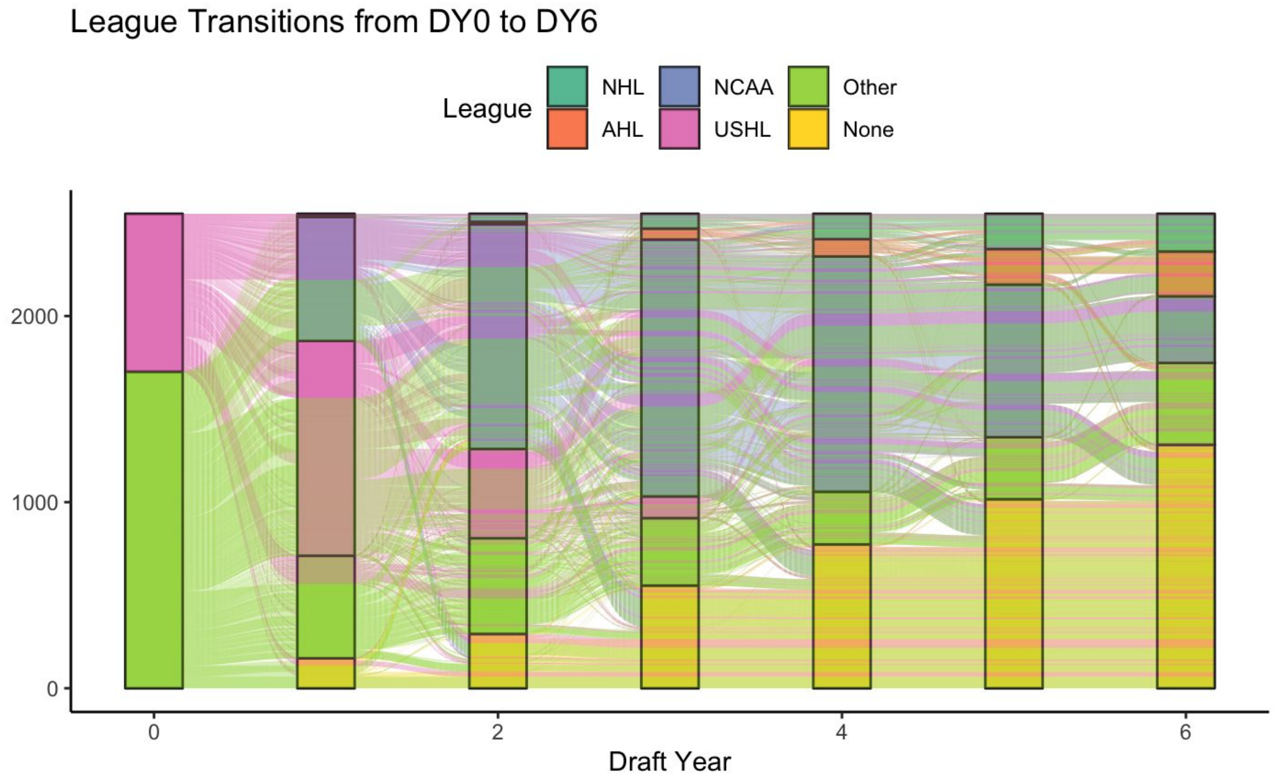

Method: Causal Inference

- Goal: determine the causal effect of development paths (Treatment **Z**) on players' future in the NHL (Response **Y**), controlling for player quality etc. (Confounders **X**)
  - Conditional Average Treatment Effect (CATE):

$$E [Y \mid Z = z1 , X] - E [Y \mid Z = z0, X]$$

  - The fundamental problem: our samples are biased

    (e.g. better prospects are more likely to enter NCAA than USHL)

Carnegie Mellon University

Method: Two solutions

- Solution 1: Control the treatment assignment mechanism; then estimate causal effect just as in a randomized experiment
  - Method: Propensity Score Matching & Weighting

- Solution 2: If we can precisely estimate a model for the outcome $Y = f(z, x) + \varepsilon$, then we can calculate CATE
  - Method: Bayesian Additive Regression Trees (BART)

Carnegie Mellon University

Method

Propensity Score Weighting

- Used to reweight the data so that we don't have any selection effect or bias in our treatment.
- We use logistic regression to predict the treatment T as well as possible from all of the predictors.
  - P(T = NCAA| X) is our **propensity score.** $\hat{e}(x) = P(T = NCAA| X)$
  - Reweight the original data based on the propensity score. We explored two weighting methods:

<u>Method 1</u>

For treatment group

$$\omega = \frac{1}{1 - \hat{e}(x)}$$

For control group

$$\omega = \frac{1}{\hat{e}(x)}$$

<u>Method 2</u>

$$\omega = \hat{e}(x)(1 - \hat{e}(x))$$

- A potential drawback of propensity scores when used for matching is that a very large number of subjects may be needed
- Propensity scores can also be used as weights in a linear model such as regression or ANOVA, so all the subjects in the control and treatment group can be used for this application

Carnegie Mellon University

Method

Propensity Score Weighting

What is the causal effect of developmental path on player success in the NHL?

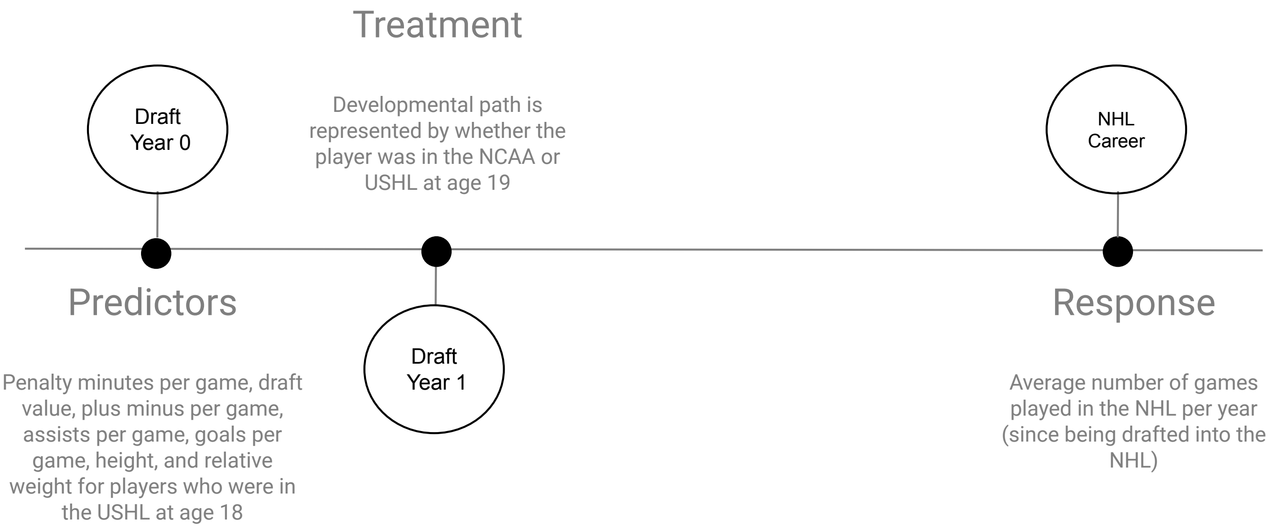

Forwards

1. Obtain the propensity score by modeling P(T = NCAA) based on predictors in draft year 0
2. Predict success in NHL with propensity score weights on the data
3. Interpret coefficient estimate for the treatment variable

Defensemen

1. Obtain the propensity score by modeling P(T = NCAA) based on predictors in draft year 0
2. Predict success in NHL with propensity score weights on the data
3. Interpret coefficient estimate for the treatment variable

Results

Propensity Score Weighting vs Matching : Forward Players

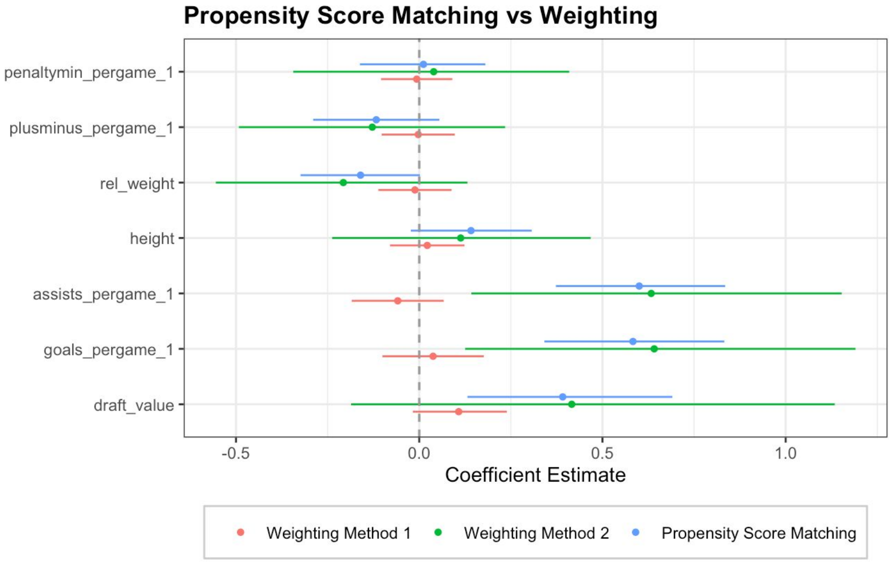

- We have **removed selection bias if the coefficient estimates are zero** when predicting treatment effect
  - This indicates that the confounders no longer have an effect on the treatment assignments
- Weighting Method 1 has 95% CIs that contain zero in all coefficient estimates
- Rel_weight is player weight after adjusting for its relationship with height

Predicting the development league at draft year 1 based on success metrics and player information in draft year 0

Carnegie Mellon University

Results

Propensity Score Weighting: Forward Players

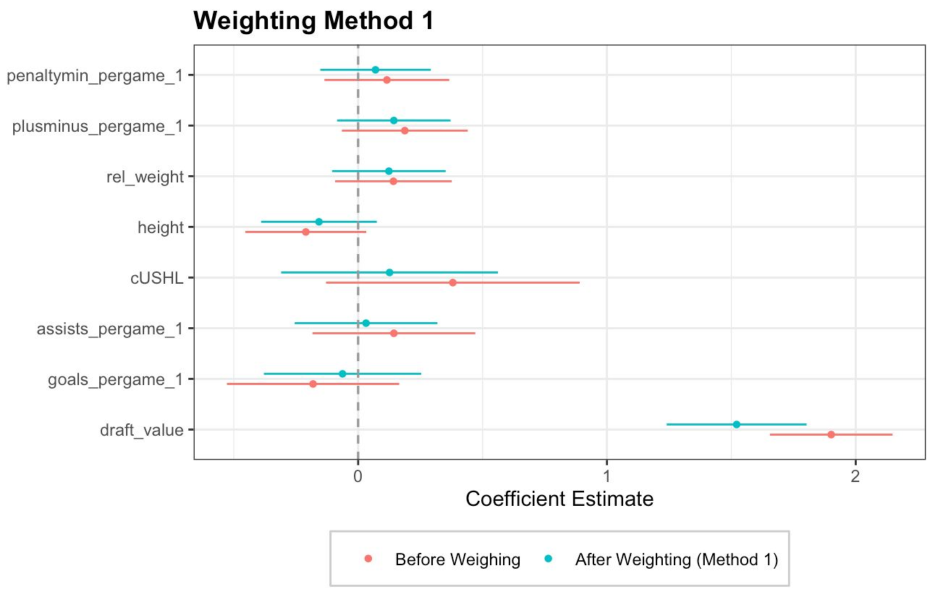

- After addressing for selection bias using weight method 1, we observe that **developmental league is not a significant predictor for success in the NHL** for forward positioned players
- Player success metrics in draft year 0 such as draft value appears to be a significant predictor after propensity score weighting
- Thus, draft value is an indicator of success in the NHL

Predicting the average number of NHL games played per year based on success metrics and player information in draft year 0 and developmental league (USHL vs NCAA) in draft year 1

Carnegie Mellon University

Results

Propensity Score Weighting vs Matching : Defensemen

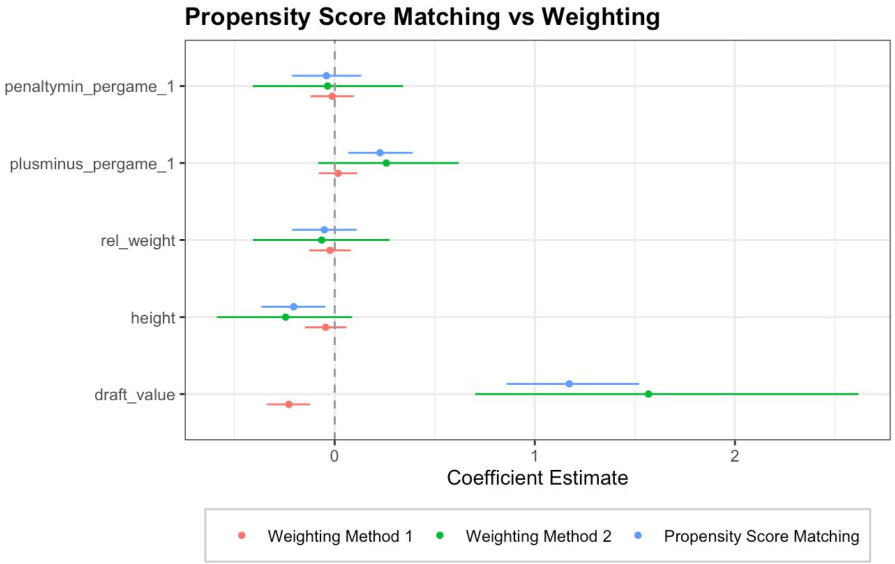

- We have **removed selection bias if the coefficient estimates are zero** when predicting treatment effect
  - This indicates that the confounders no longer have an effect on the treatment assignments
- Weighting Method 1 has 95% CIs that contain zero in most coefficient estimates
- Removed goals and assists per game due to correlation with draft value and lower adjusted r squared
- Rel_weight is player weight after adjusting for its relationship with height

Predicting the development league at draft year 1 based on success metrics and player information in draft year 0

Results

Propensity Score Matching: Defensemen

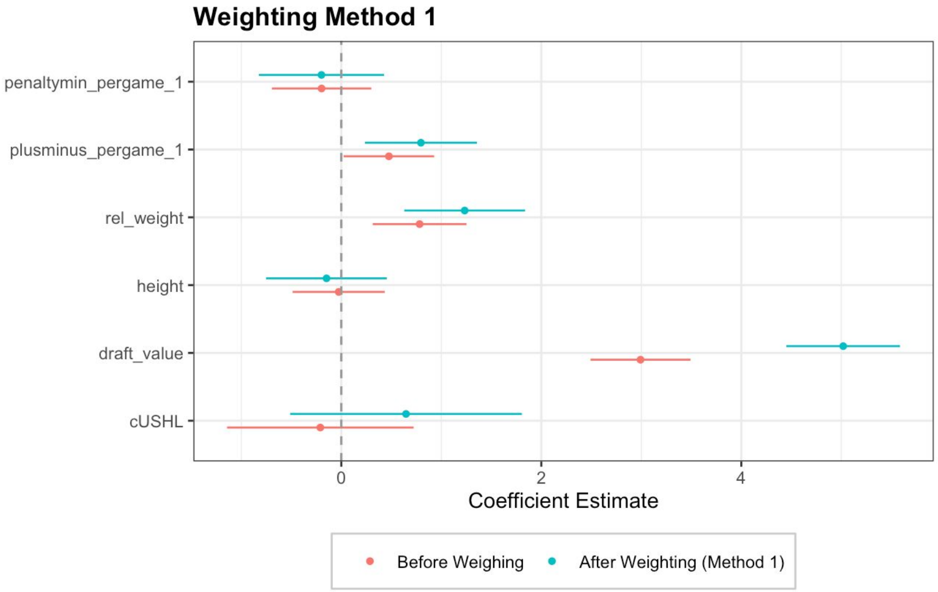

- After addressing for selection bias using weight method 1, we observe that **developmental league is not a significant predictor for success in the NHL** for defensemen
- Player success metrics in draft year 0 such as draft value appears to be a significant predictor after propensity score weighting
- Thus, draft value is an indicator of success in the NHL

Predicting the average number of NHL games played per year based on success metrics and player information in draft year 0 and developmental league (USHL vs NCAA) in draft year 1

Carnegie Mellon University

Method

Bayesian Additive Regression Trees (BART)

- Estimate Y = f (z, x) + ε, where ε ~ N(0, σ^2) using a **sum-of-trees** model

- The idea is to fit a bunch of **weak-learning** (small) trees
  each fitting to the residuals of the previous trees
  additively combine these trees to reduce bias, similar to boosting.

- Introduce a **regularization prior** to avoid overfitting.
  controls the size of the trees (T), the magnitude of the outputs trees (M), and the value of σ^2.

- Compute the **posterior** using **Markov Chain Monte Carlo** (MCMC)
- At each iteration of MCMC, T, M and σ are redrawn to seek a good posterior.

- Using BART, we can calculate CATE by $\frac{1}{n}\sum_{i=1}^{n} f(1, x_i) - f(0, x_i)$

Carnegie Mellon University

Method

Bayesian Additive Regression Trees (BART)

- Selected 2826 players (1830 Forward and 996 Defensemen) who played in USHL in their initial draft year
  - 233 players who played in either USHL or NCAA in the following year

- Predictors:
  - Position (Forward/Defensemen), Height, Weight, Penalty minutes, Draft value, League in the following season

- Response variable:
  - The average number of games played in the NHL per season

Result -- Treatment Effects (NCAA vs. USHL)

BART

Y = The average number of games played in the NHL per season

Observed Average Difference = average Y for the NCAA group - average Y for the USHL group

$$CATE = E[Y \mid Z = NCAA, X] - E[Y \mid Z = USHL, X]$$

	Observed Average Difference	CATE using BART	CATE from propensity score weighting
Forward	13.36513	0.0541	2.0297
Defense*	14.65782	0.0554	0.1871

*Defense: "D/F", "D/LW", "D"

Carnegie Mellon University

Result -- Variable Importance

BART

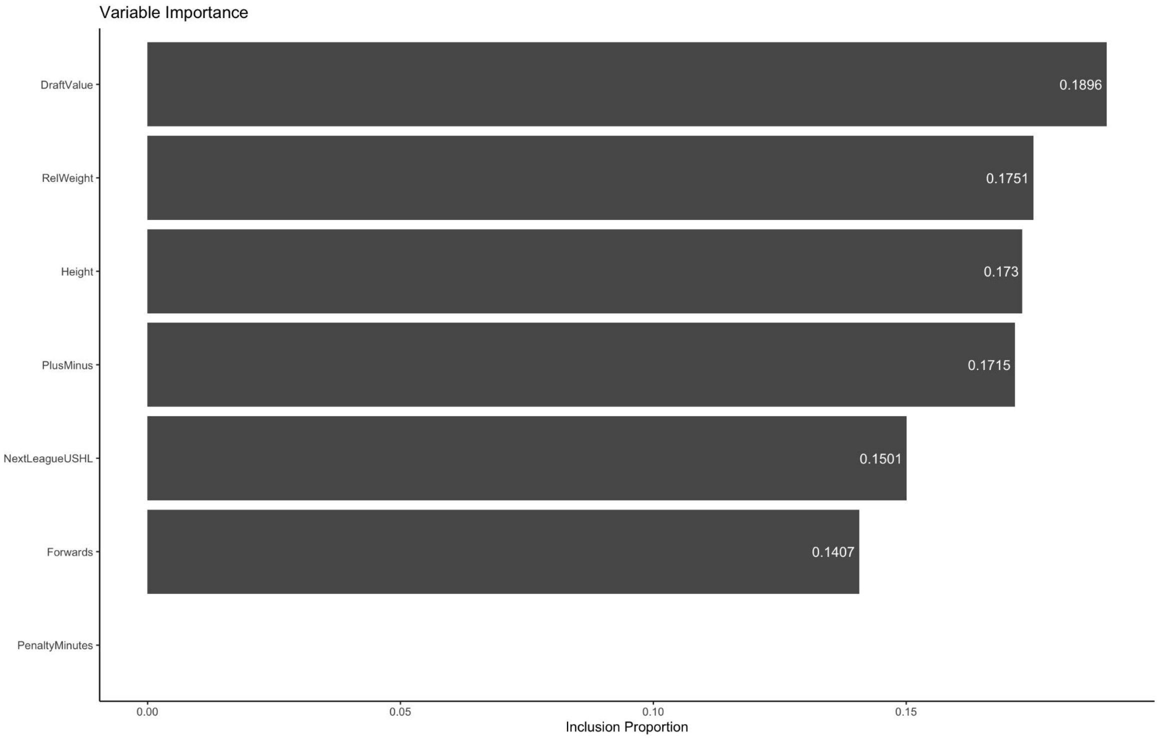

NextLeagueUSHL = binary
(1 if USHL, 0 if NCAA)

Forwards = binary
(1 if Forwards, 0 if Defensemen)

Developmental path and position take less impact in predicting for the average number of NHL games per season

Penalty minutes seems irrelevant

Carnegie Mellon University

Result -- Variable Importance

BART

Variable Importance, separated by position

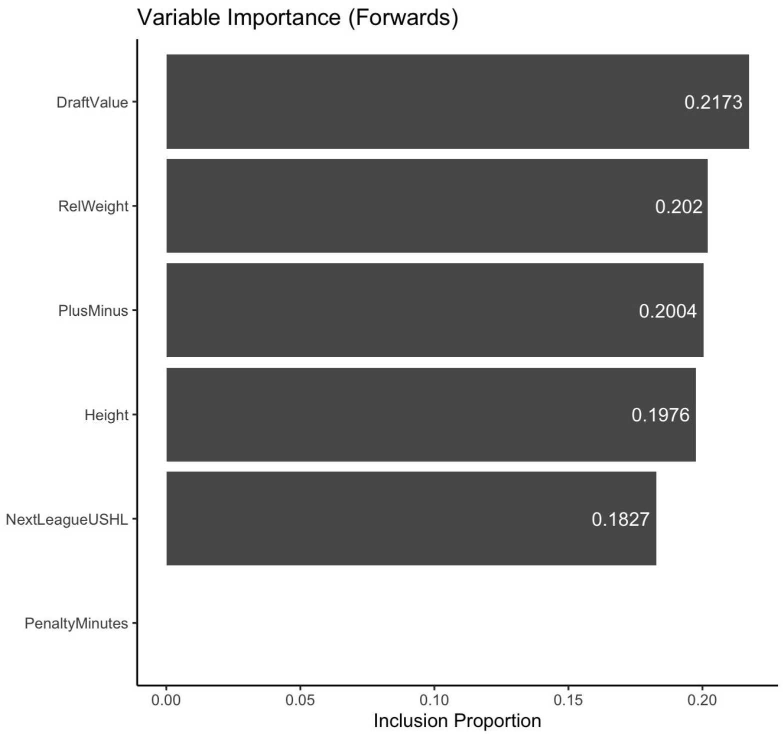

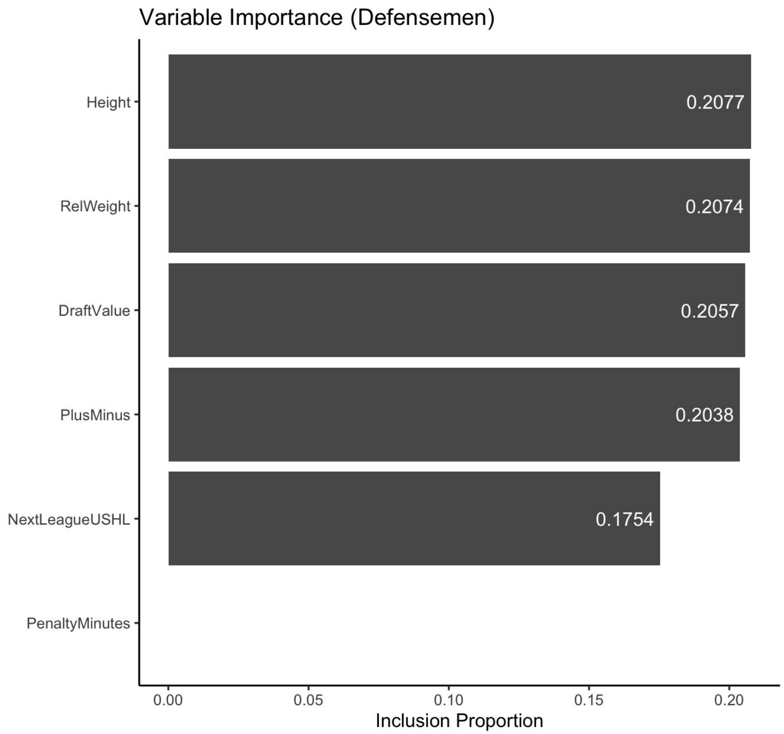

What if …

I have one more goal?

I have one more assist?

I am 1 inch taller?

I am 1 pound heavier?

Result

BART

- What if one more goal for forward players?

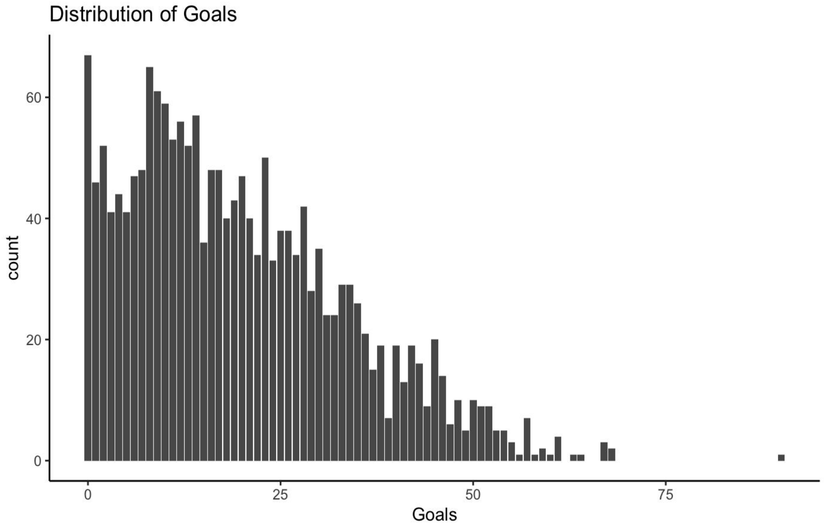

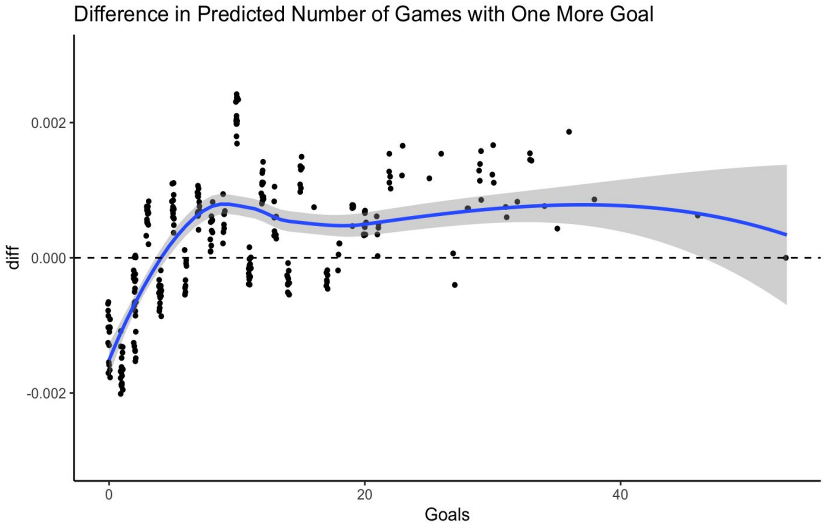

Carnegie Mellon University

Result

BART

- What if one more goal for defense players?

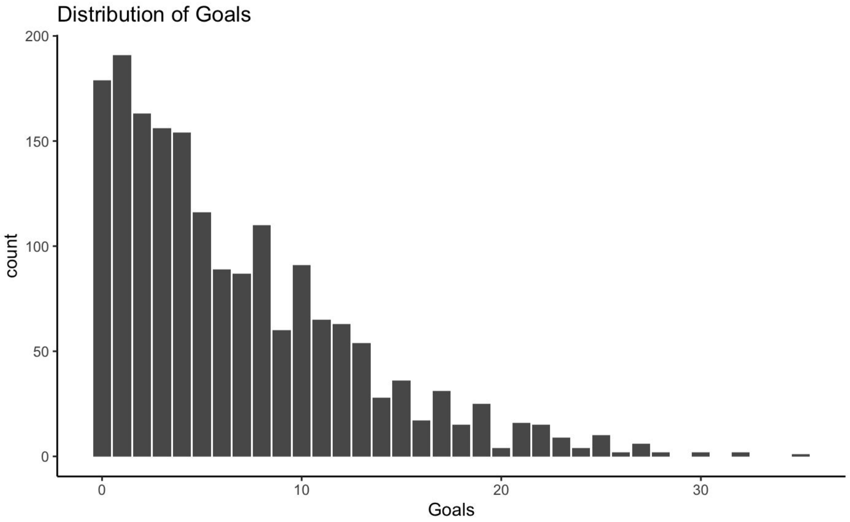

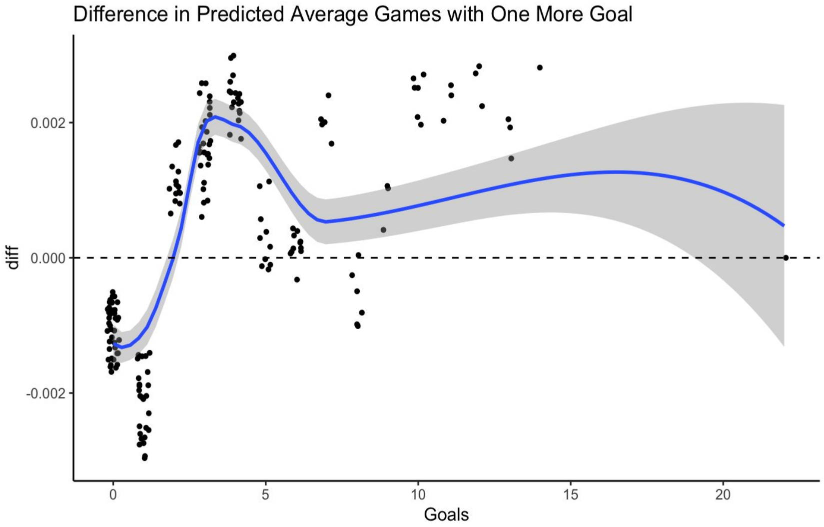

Carnegie Mellon University

Result

BART

- What if one more assist for forward players?

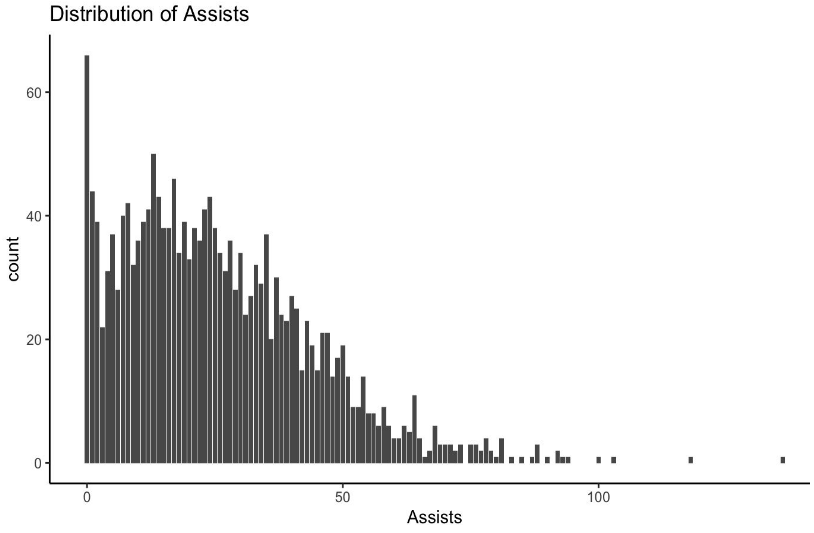

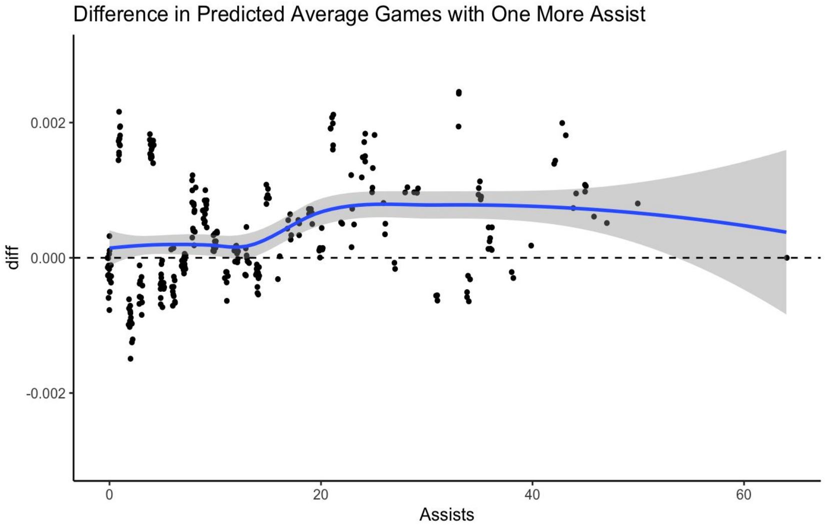

Result

BART

- What if one more assist for defense players?

Distribution of Assists

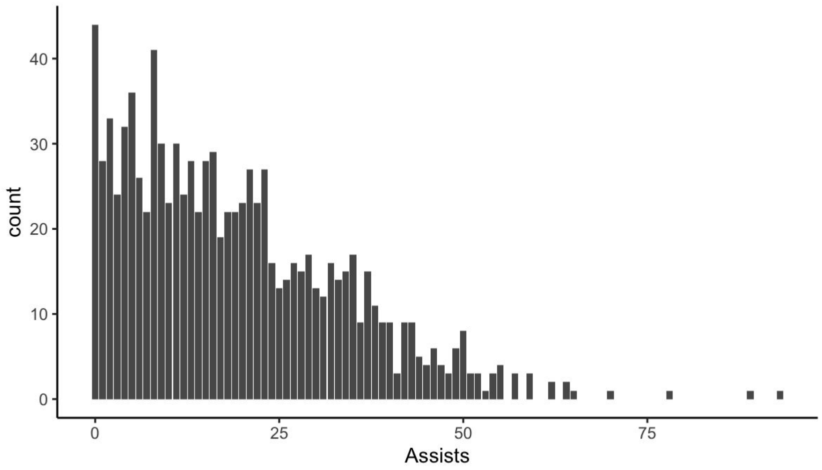

Difference in Predicted Average Games with One More Assist

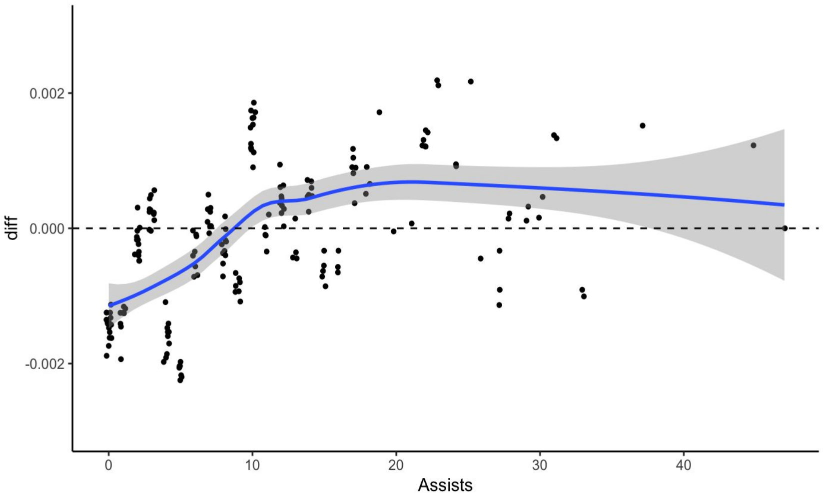

Result

BART

- What if one inch taller or one pound heavier?

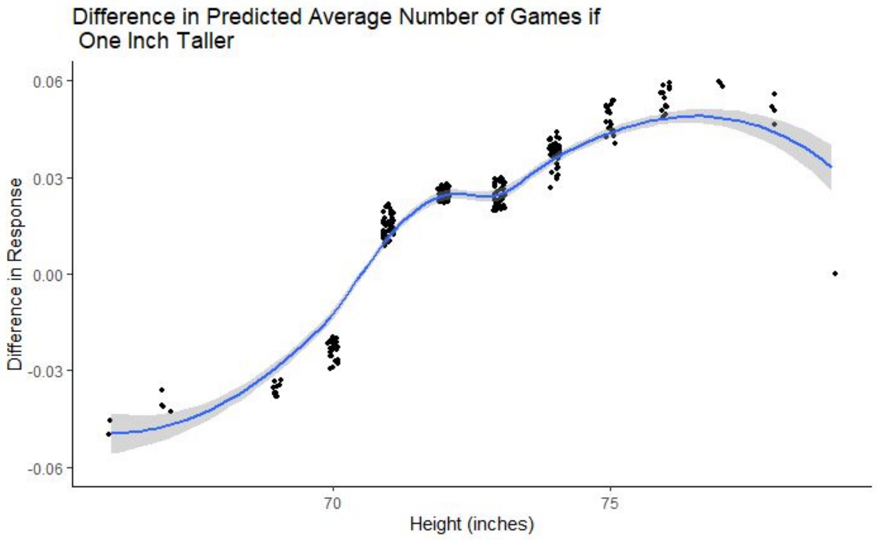

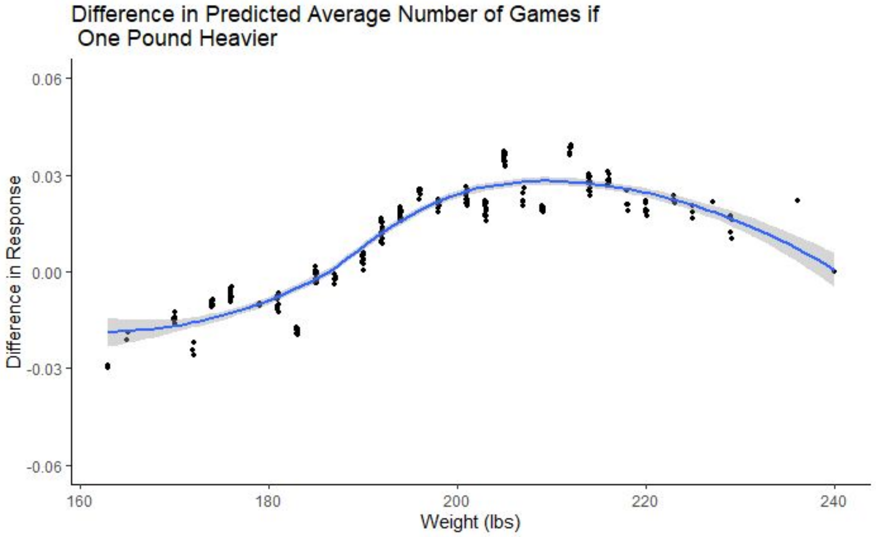

Carnegie Mellon University

Discussion

- Developmental path does not seem to have much causal effect

- We only took players just entering their draft year
    - Potential causal effect may change (may have more impact) later in their developmental phase

- We only considered two main developmental leagues (USHL and NCAA)
    - So much other developmental leagues, including overseas leagues

Q&A

Thank You!

Parameter Testing with Priors

The prior for the BART model has three components:

(1) the tree structure itself,

(2) the leaf parameters given the tree structure

(3) the error variance σ^2 which is independent of the tree structure and leaf parameters

α:

□:

K: larger k means more model regularization

μ:

q:

Carnegie Mellon University

33

Parameter Testing with Priors

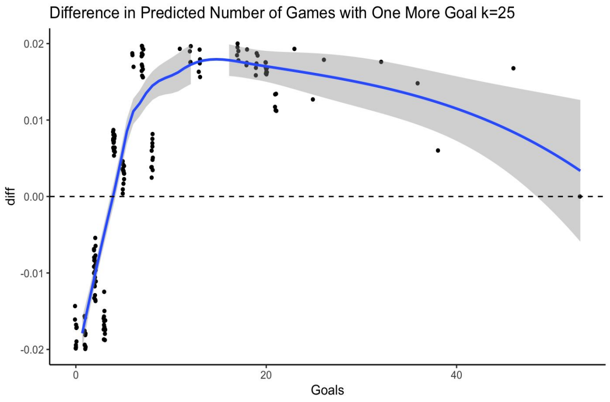

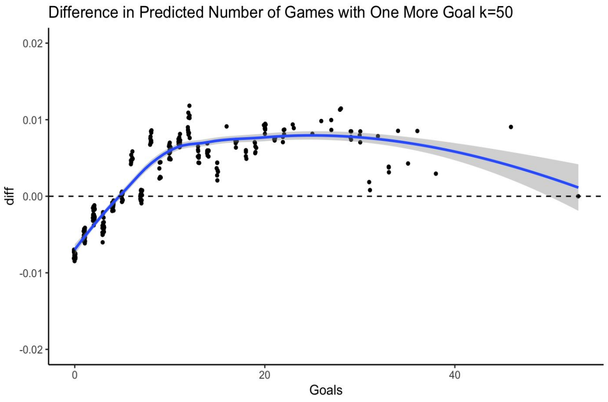

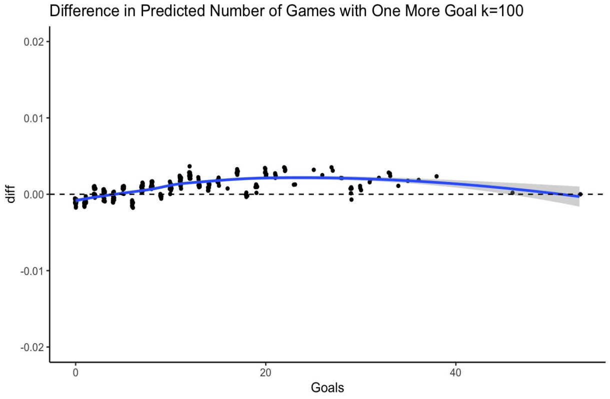

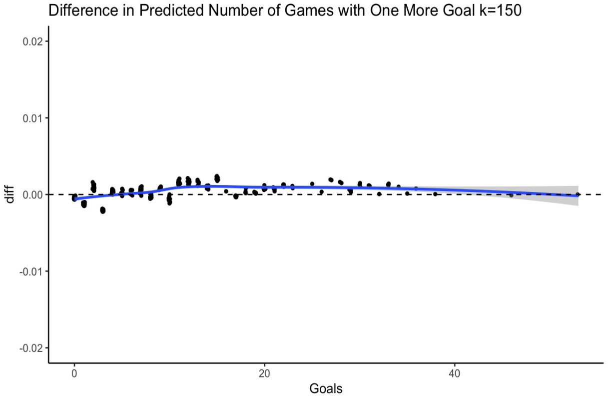

Result

BART

- What if one inch taller for defend players?

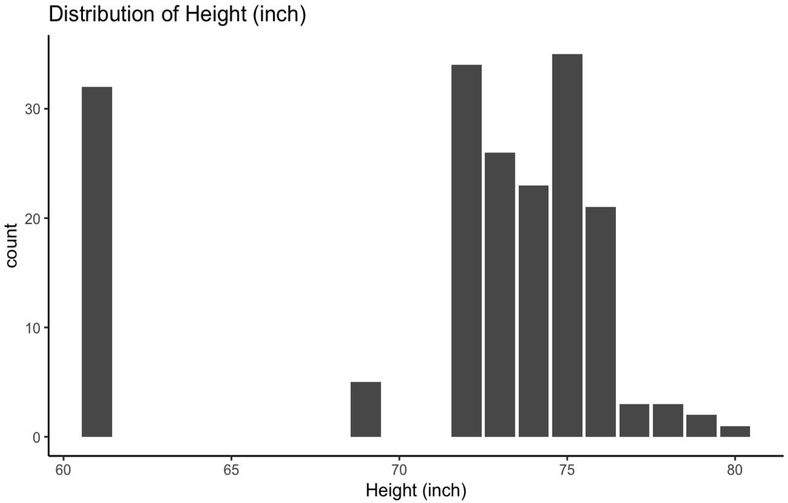

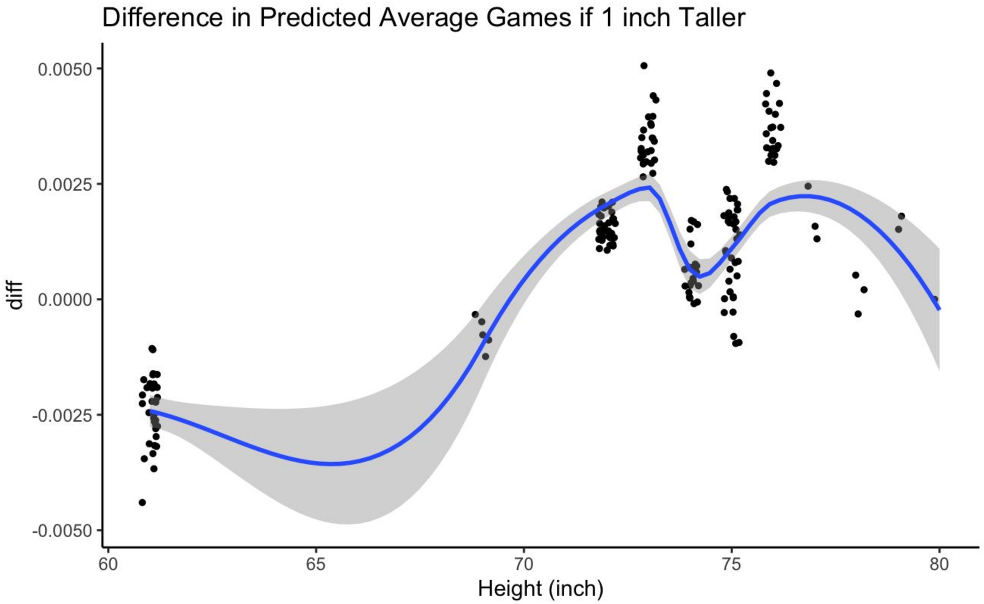

Result

BART

- What if one centimeter/(0.39 inches) taller for forward players?

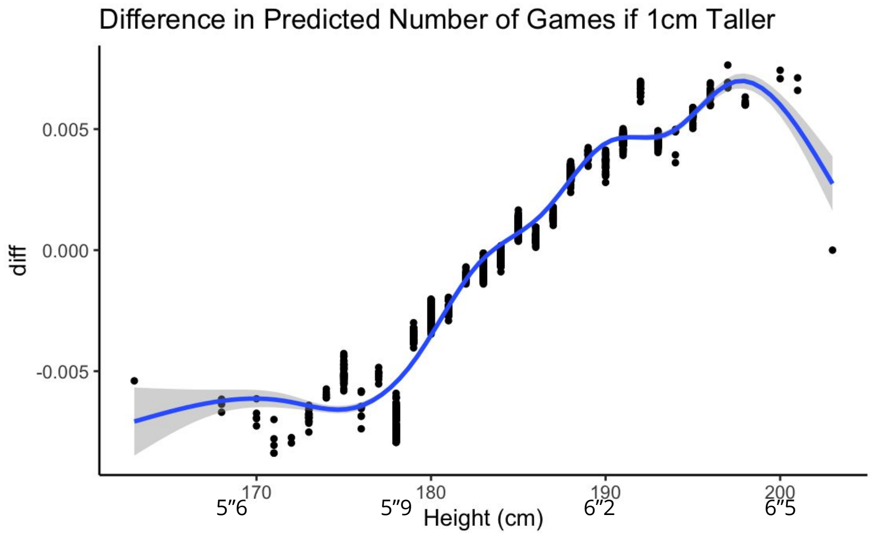

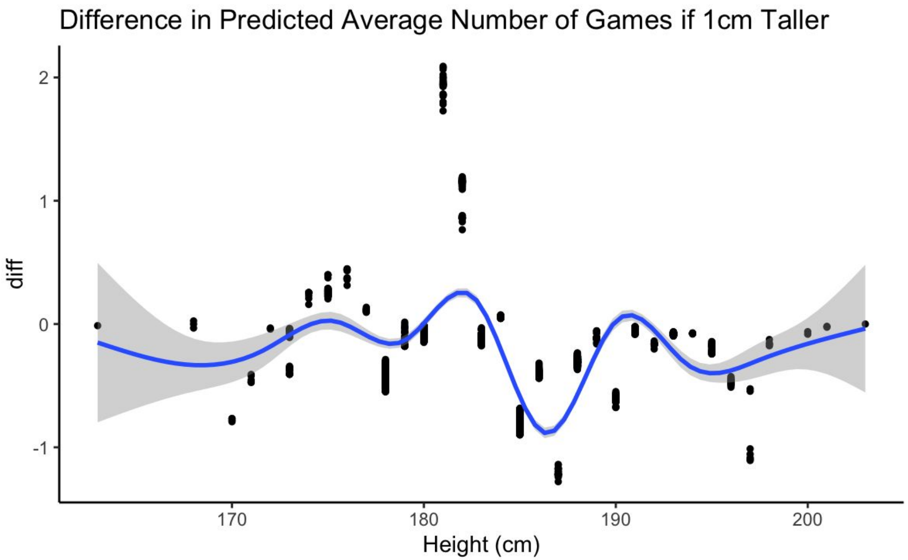

Carnegie Mellon University

Result

BART

- What if one centimeter/(0.39 inches) taller for defense players?

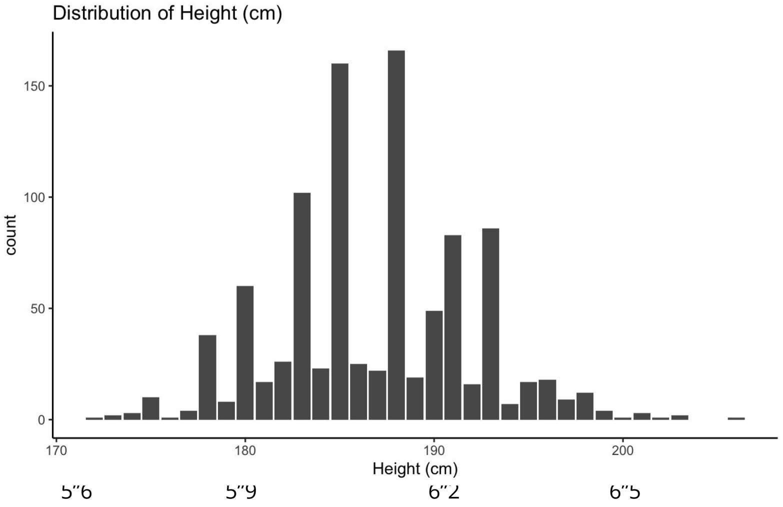

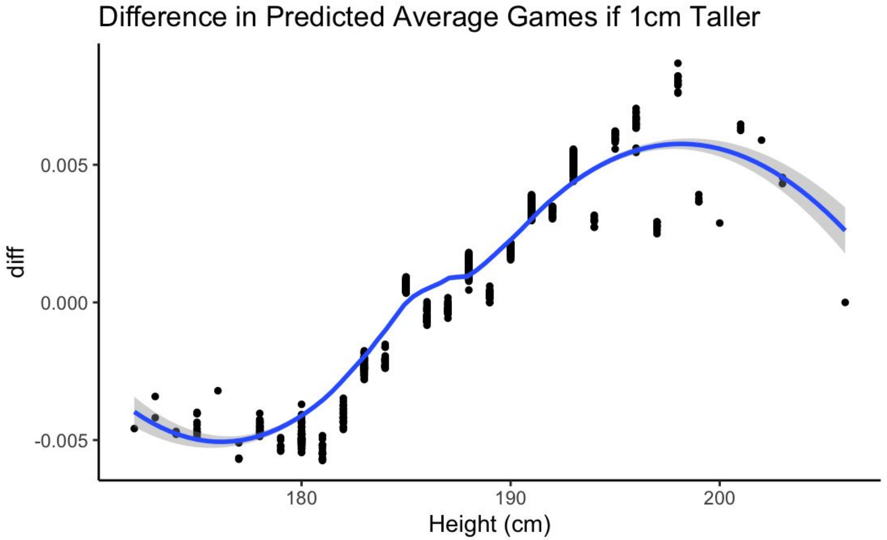

Result

BART

- What if 1kg/2.2 pounds heavier for forward players?

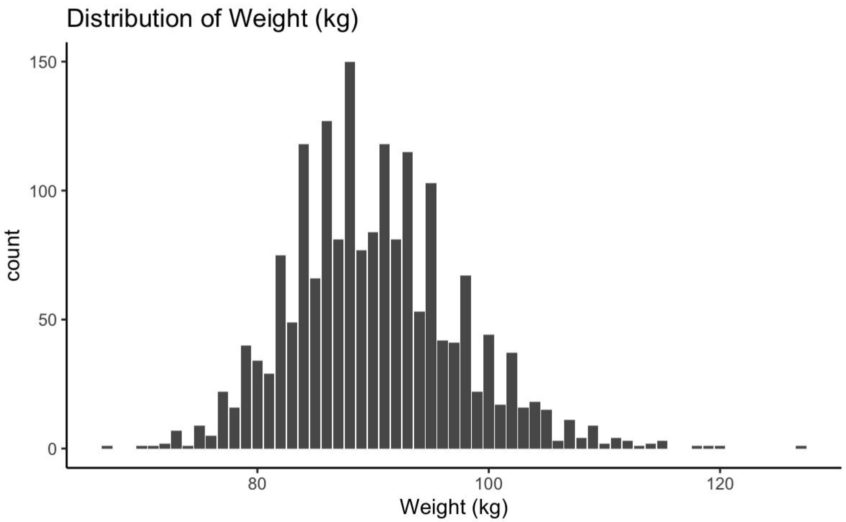

176.38 200 264 pounds

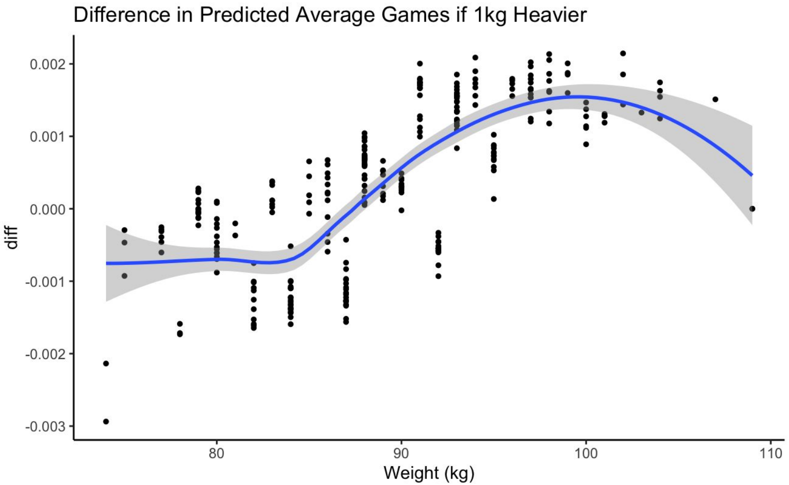

Result

BART

- What if one 1kg/2.2 pounds for defense players?

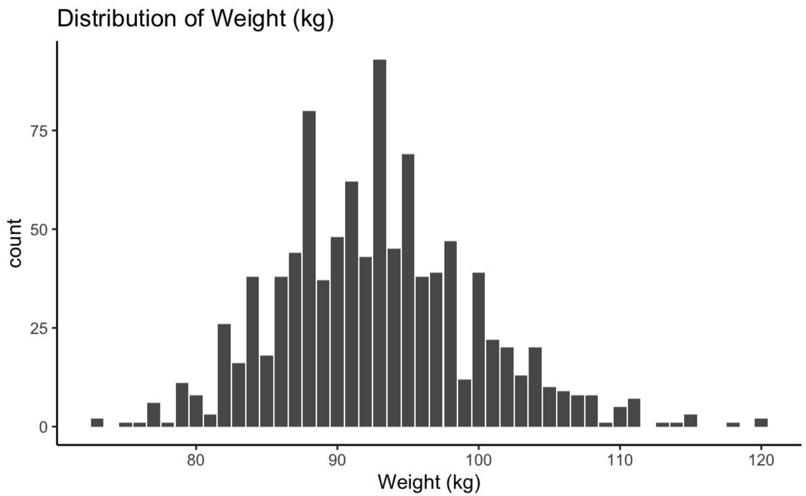

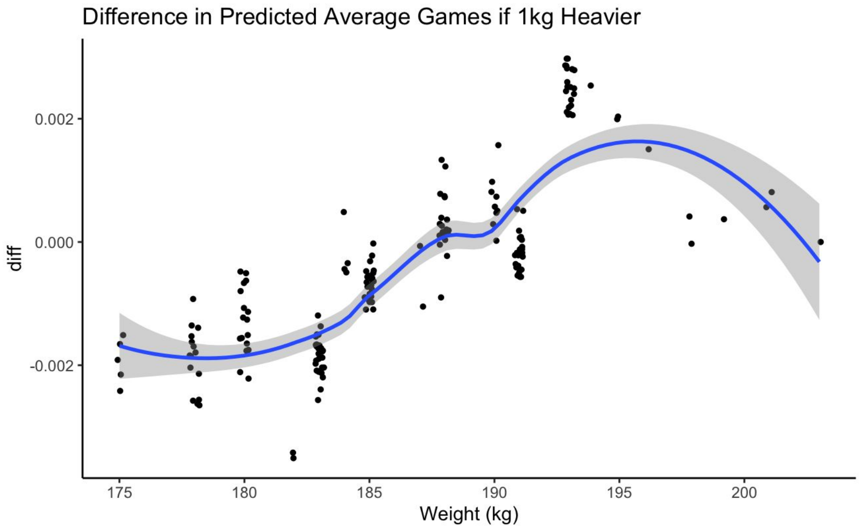

Carnegie Mellon University