Causal Inference and Machine Learning for Social Science

Jennie E. Brand

Department of Sociology

Department of Statistics and Data Science

California Center for Population Research (CCPR)

Center for Social Statistics (CSS)

UCLA

Talk Prepared for Future Directions for Social and Behavioral Science Methodologies in the Next Decade

September 2024

The potential outcomes framework offers a conceptual apparatus for defining causal effects

The potential outcomes framework offers a conceptual apparatus for defining causal effects

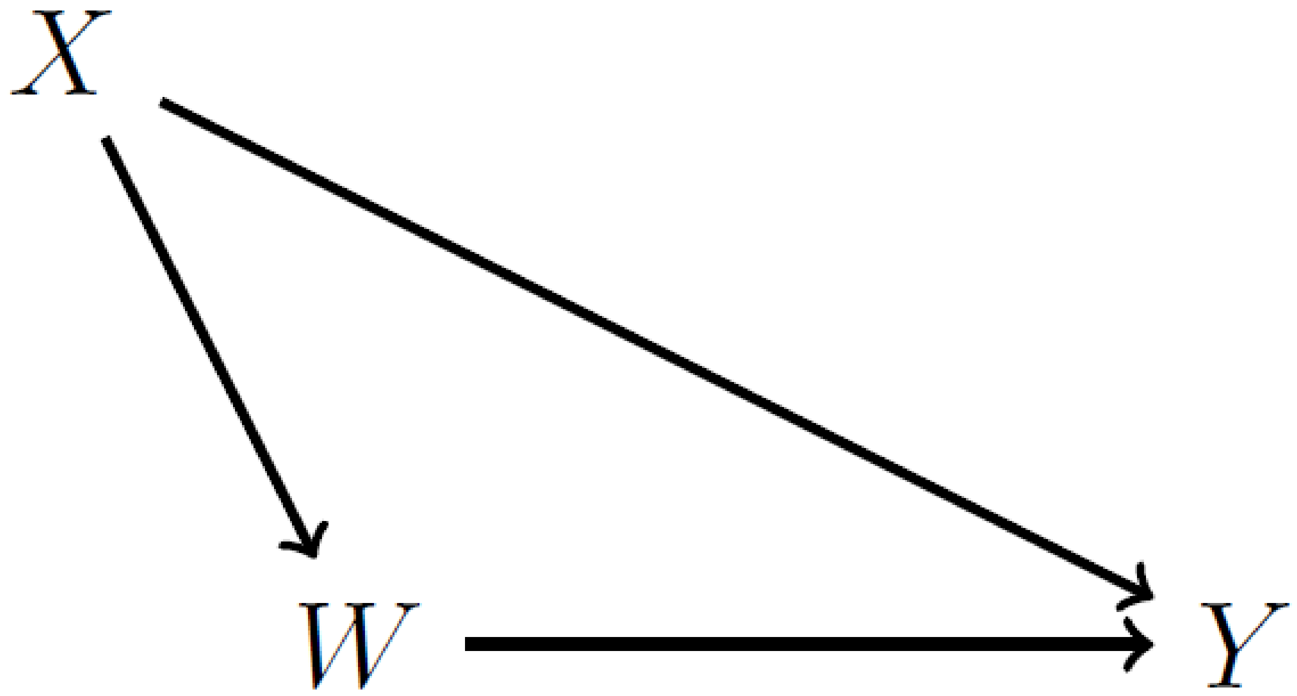

Direct Acyclic Graph under Unconfoundedness

Note: W denotes treatment status, Y denotes the outcome of interest, and X denotes observed pretreatment confounders.

Researchers draw on various methods to estimate causal effects

- To adjust for pretreatment heterogeneity, researchers may use **regression imputation**, propensity score matching (**PSM**), and inverse probability weighting (**IPW**)

Researchers draw on various methods to estimate causal effects

- To adjust for pretreatment heterogeneity, researchers may use **regression imputation**, propensity score matching (**PSM**), and inverse probability weighting (**IPW**)
- **Flexible machine learning** methods can be used to fit the outcome or propensity score models

Researchers draw on various methods to estimate causal effects

- To adjust for pretreatment heterogeneity, researchers may use **regression imputation**, propensity score matching (**PSM**), and inverse probability weighting (**IPW**)
- **Flexible machine learning** methods can be used to fit the outcome or propensity score models
  - Researchers have **adapted machine learning** methods **to estimate causal parameters** to mitigate concerns central to causal inference
  - Machine learning methods perform well when combined with the so-called "**doubly robust estimators**" of average treatment effects
  - To minimize overfitting, we use sample splitting or cross-fitting

Researchers should routinely attend to response variation, i.e., ‘treatment effect heterogeneity’

- Individuals differ not only in pre-treatment characteristics (i.e., **pre-treatment heterogeneity**), but also in how they respond to a common treatment (i.e., **treatment effect heterogeneity**)
- Treatment effect heterogeneity has important implications for **social and behavioral research and social policy**

Researchers should routinely attend to response variation, i.e., ‘treatment effect heterogeneity’

- Individuals differ not only in pre-treatment characteristics (i.e., **pre-treatment heterogeneity**), but also in how they respond to a common treatment (i.e., **treatment effect heterogeneity**)
- Treatment effect heterogeneity has important implications for **social and behavioral research and social policy**
  - Yields important insights into how scarce social resources are distributed in an unequal society
  - Helps extrapolate findings to diverse populations and contexts
  - Plays a critical role in guiding evidence-based policy

Researchers should routinely attend to response variation, i.e., ‘treatment effect heterogeneity’

- Individuals differ not only in pre-treatment characteristics (i.e., **pre-treatment heterogeneity**), but also in how they respond to a common treatment (i.e., **treatment effect heterogeneity**)
- Treatment effect heterogeneity has important implications for **social and behavioral research and social policy**
  - Yields important insights into how scarce social resources are distributed in an unequal society
  - Helps extrapolate findings to diverse populations and contexts
  - Plays a critical role in guiding evidence-based policy
- Various methods help us **uncover response variation**, including machine learning

Effects of Completing College on Low-Wage Work

Research objective

- Examine the distribution of effects of college completion on low-wage work

Data and variables

- National Longitudinal Survey of Youth 1979 (NLSY)
- Treatment: College completion by age 25
- Outcome: Low-wage work age 25-50

Brand, Jennie E., Jiahui Xu, Bernard Koch, and Pablo Geraldo. 2021. "Uncovering Sociological Effect Heterogeneity using Tree-Based Machine Learning." *Sociological Methodology* 51(2):189-223.

Brand, Jennie E. 2023. *Overcoming the Odds. The Benefits of Completing College for Unlikely Graduates*. New York: Russell Sage Foundation.

Propensity Score Specification

Step 1: Baseline set of covariates K_B

Step 2: Consider $K - K_B$ additional possible covariates in turn

- 176 logistic regressions, resulting in a model with 22 K_L linear terms

Step 3: Consider all possible higher order and interaction terms $[K_L (K_L + 1)/2]$ in turn

- 253 additional terms, 3,527 regressions, resulting in a model with 1 higher order term and 12 interaction terms

Propensity Score Specification

Sociodemographic

Race

Sex

Residence

Family Background

Parents' income

Parents' education

Fathers' occupation

Family structure

School characteristics

School disadvantage

Cognitive and psychosocial

Test scores (ASVAB)

College-prep. program

Locus of control

Delinquency

Expectations / aspirations

Friends' aspirations

Family formation

Martial status

Had a child

Heterogeneous Effects of College on Low-Wage Work

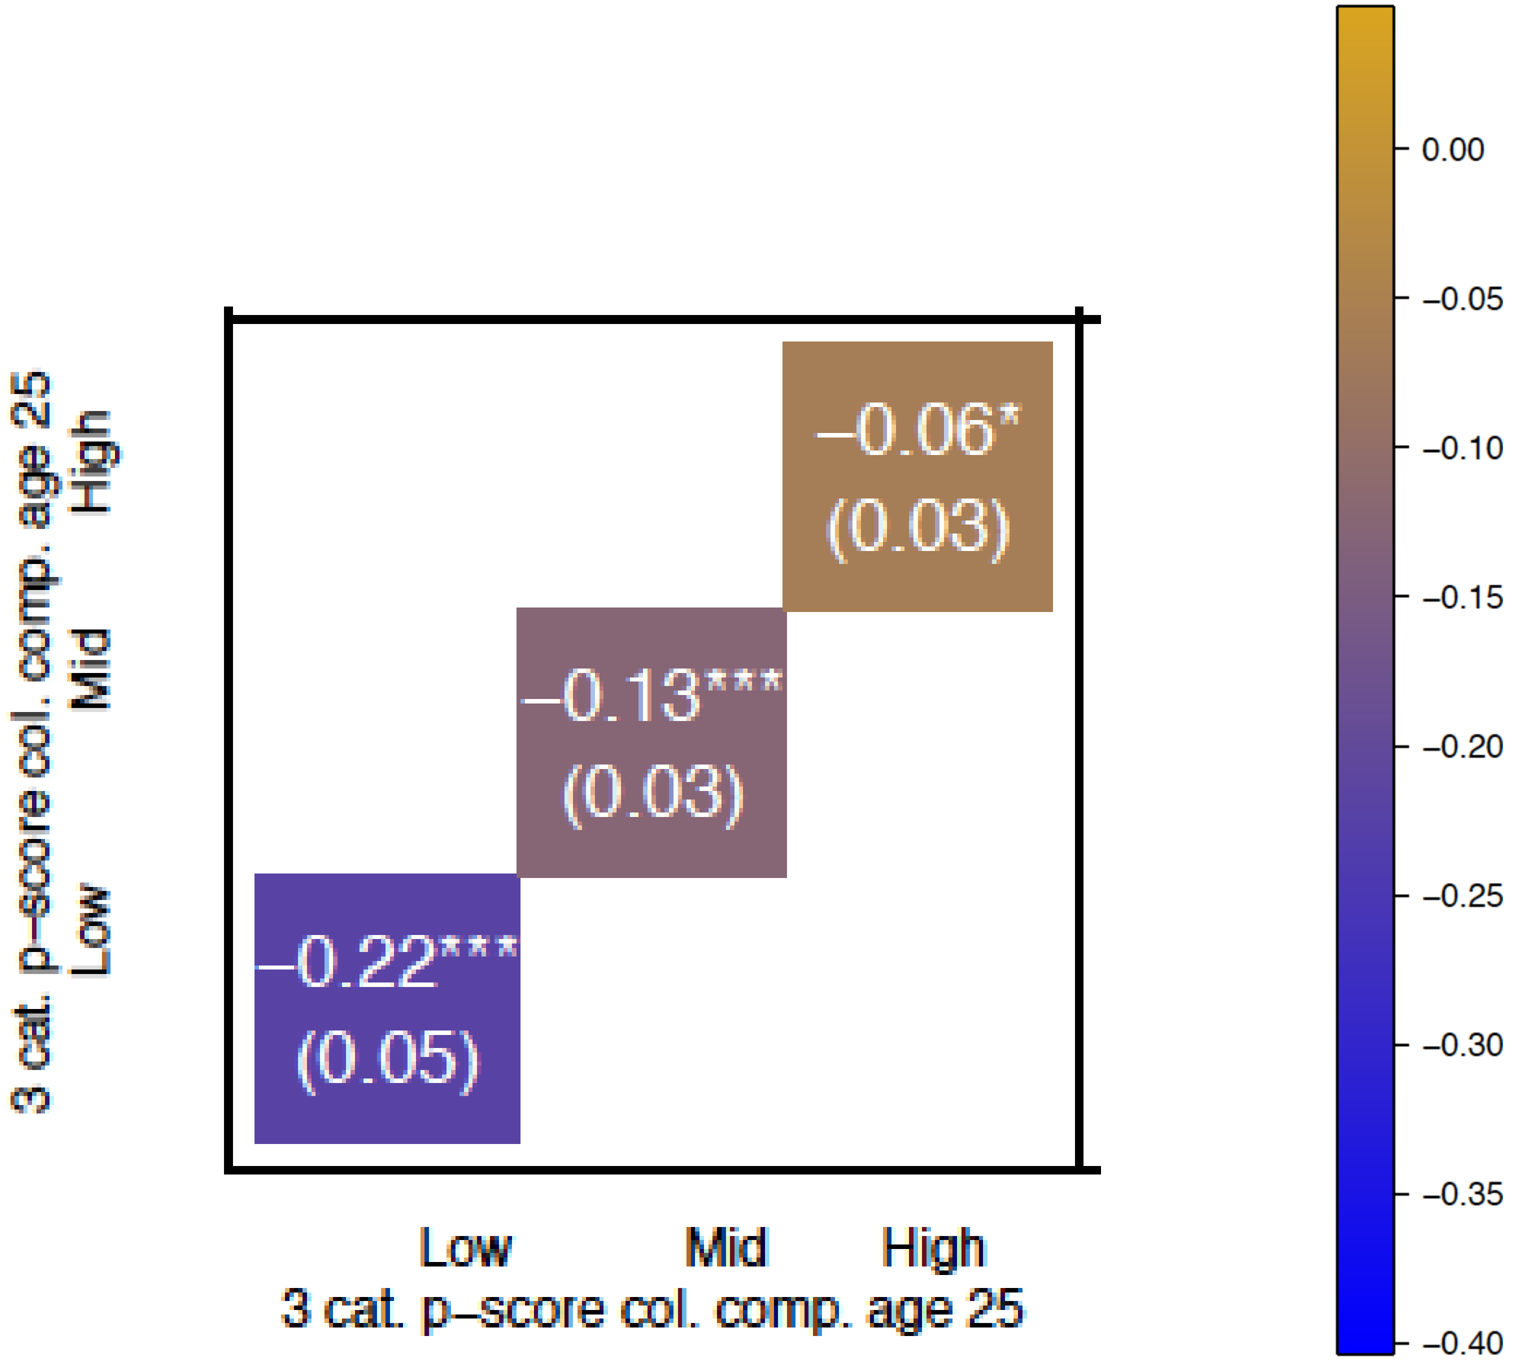

Heterogeneous Effects of College on Low-Wage Work

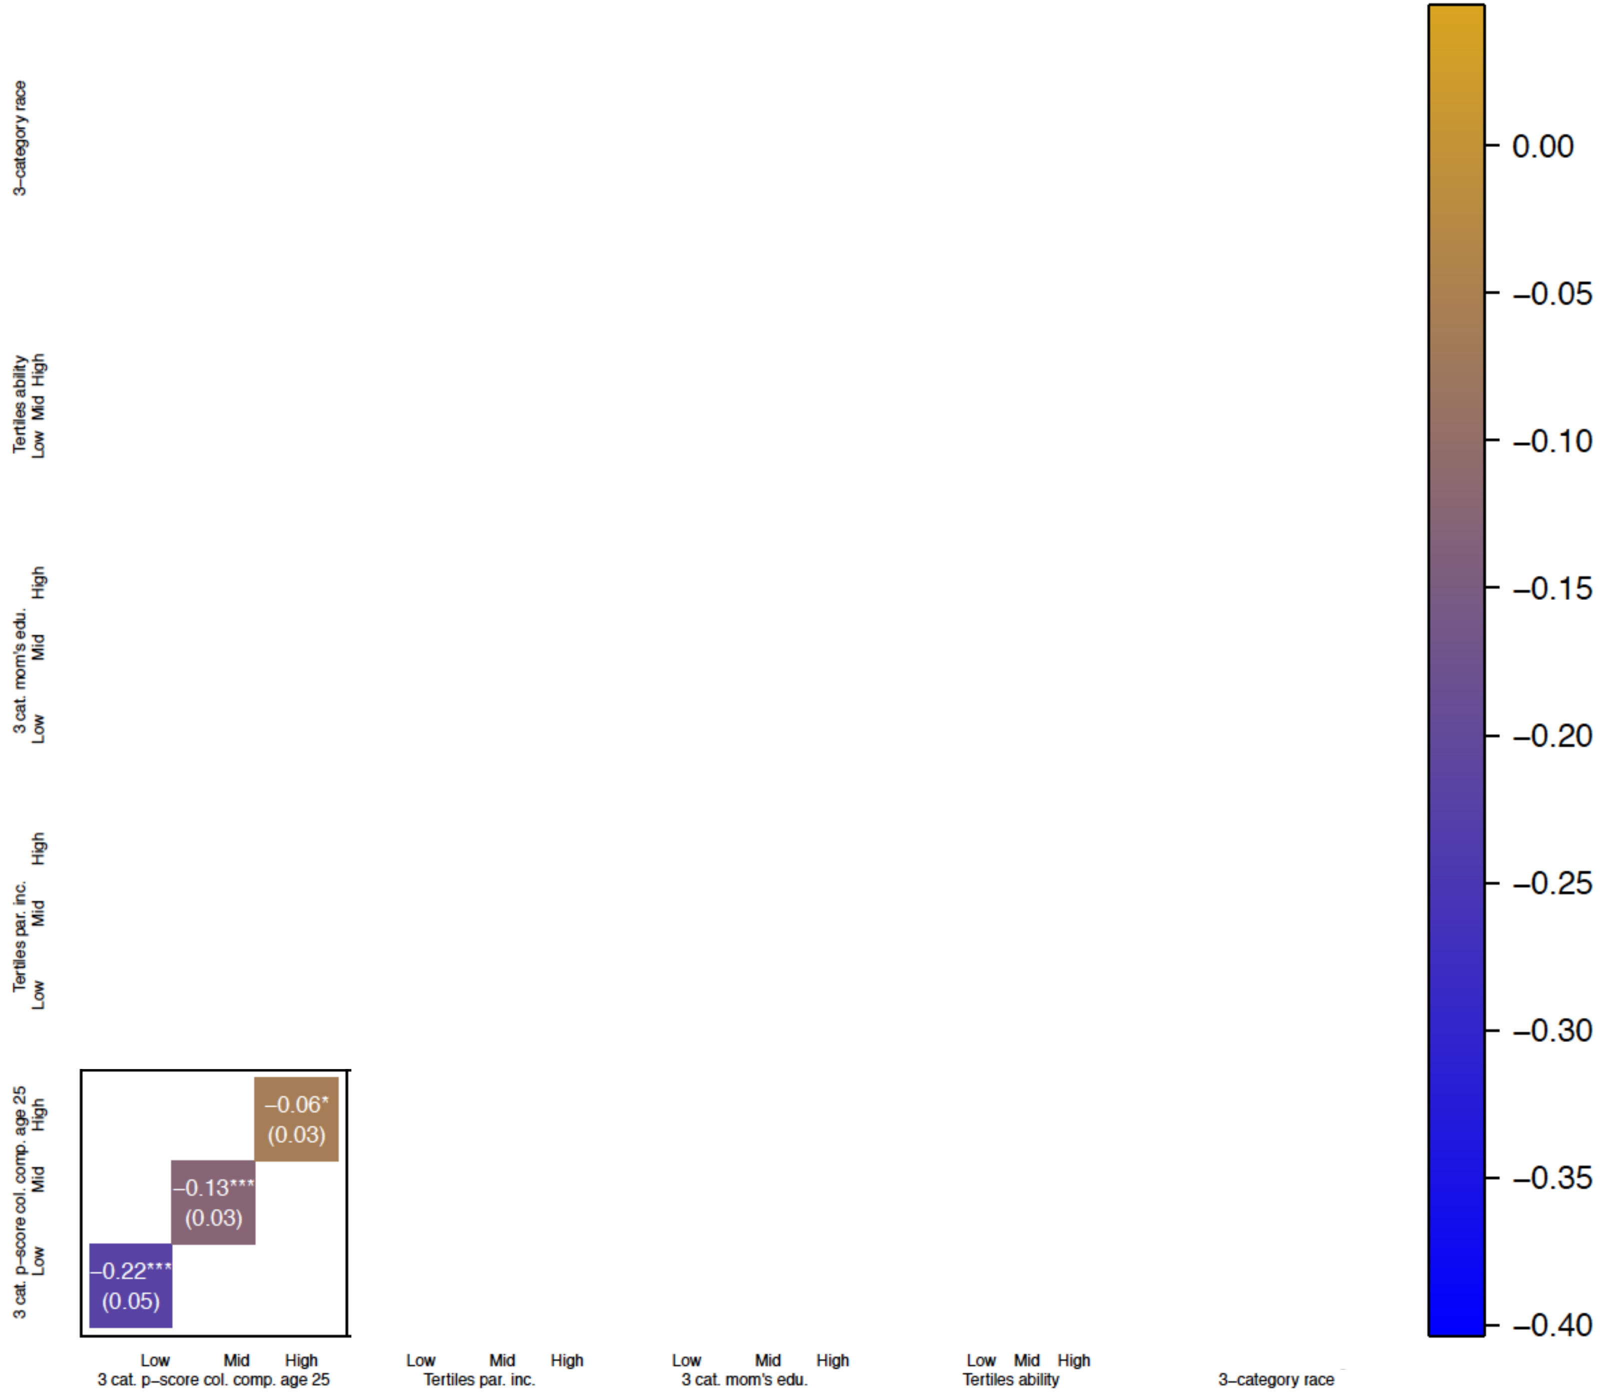

Heterogeneous Effects of College on Low-Wage Work

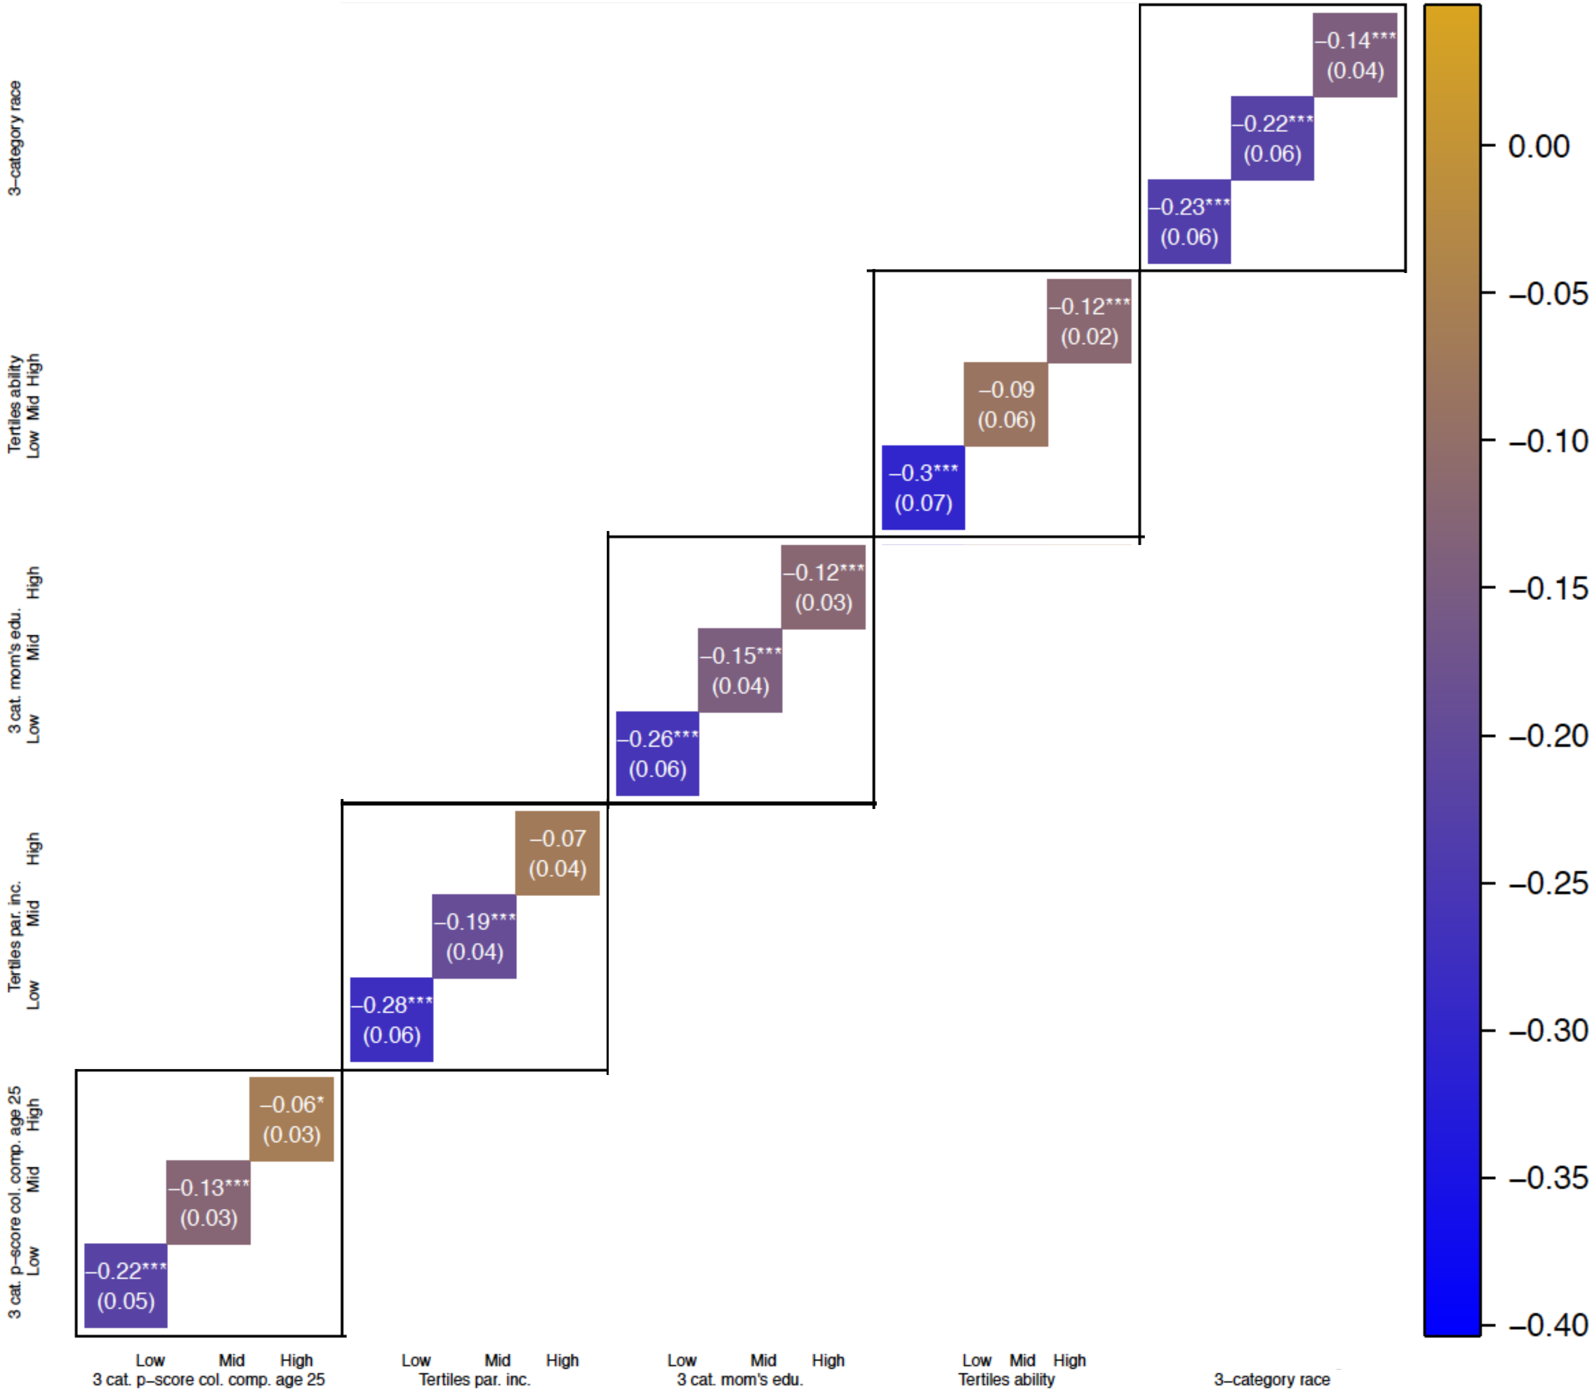

Heterogeneous Effects of College on Low-Wage Work

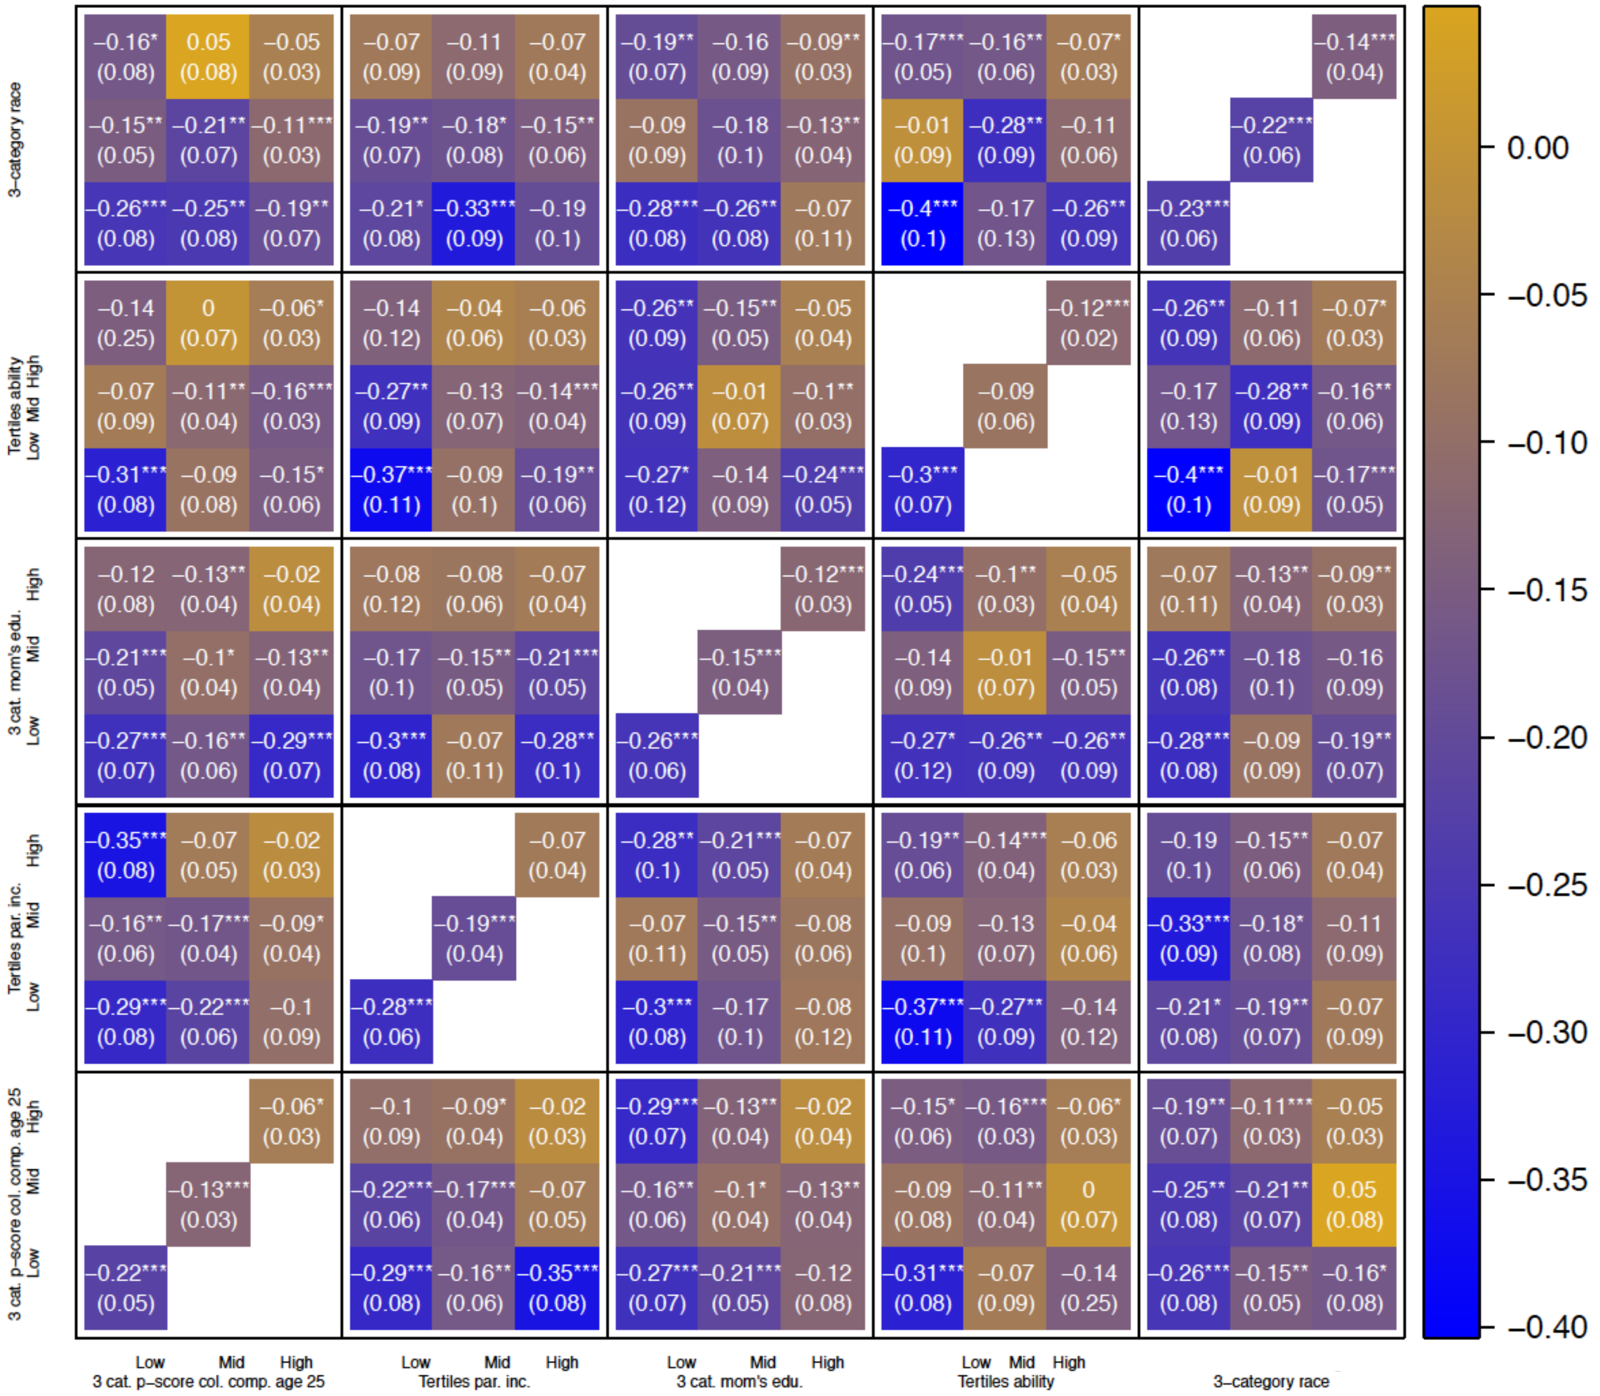

Are there important sources of variation that researchers may not have considered prior to data analysis?

- **Researchers routinely explore their data** to determine if subgroups show meaningful differences in effect estimates
- If researchers select which interactions to report from these analyses and do not draw on **cross-validation procedures** or **multiple-testing penalties**, they are subject to **incorrectly failing to accept the null hypothesis**
- It may be unclear which of the large number of possible **joint covariates and thresholds** are **best to consider before analyses**
- Statisticians and social and computer scientists have recently made progress in **merging machine learning methods and causal inference**
- **Decision trees** uncover new sources of variation

Decision trees recursively partition the data by covariates into increasingly smaller subsets

- Covariates and thresholds are selected that minimize the in-sample loss function, and the sample is split into two new partitions

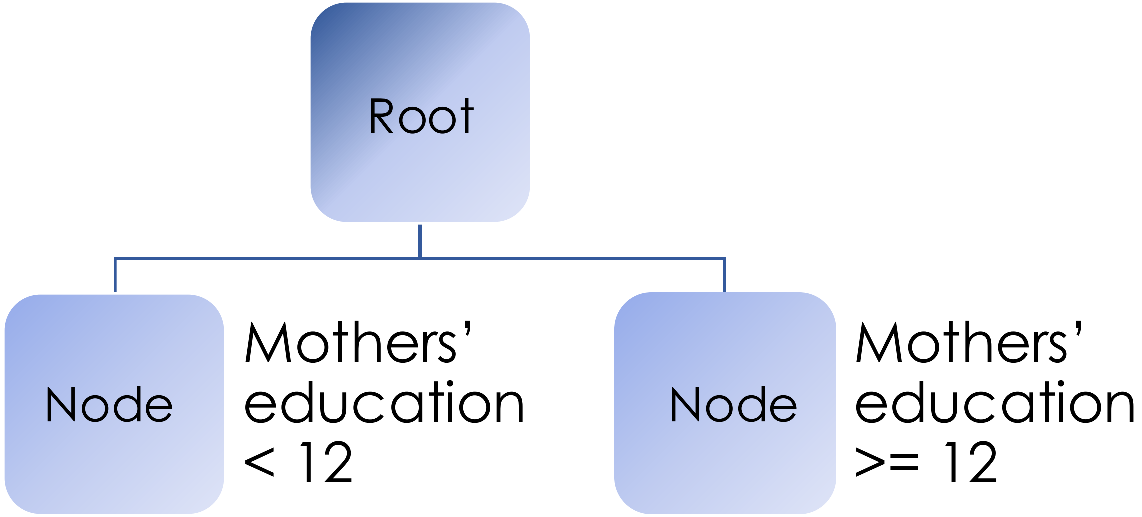

Decision trees recursively partition the data by covariates into increasingly smaller subsets

- Covariates and thresholds are selected that minimize the in-sample loss function, and the sample is split into two new partitions

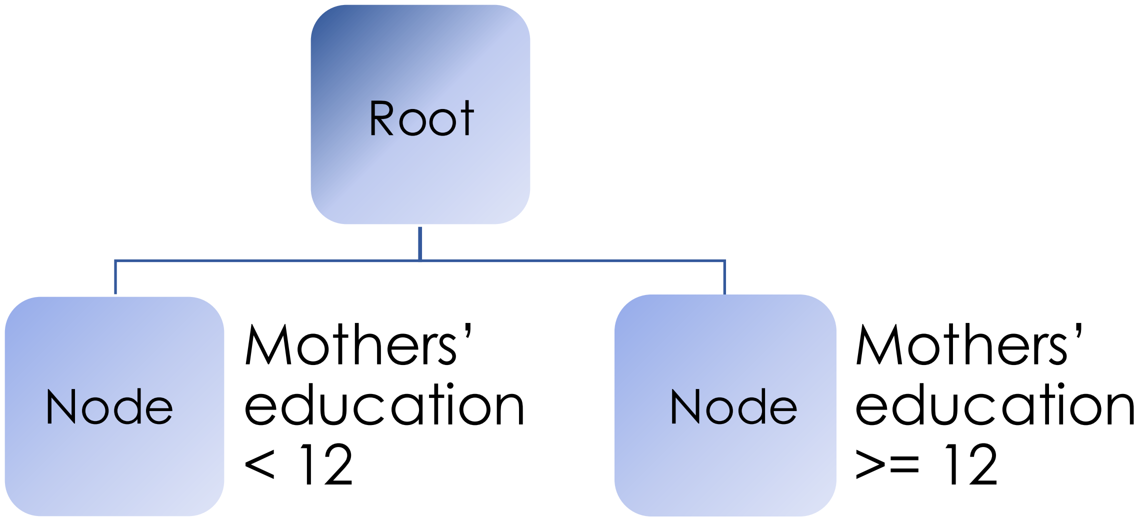

Decision trees → Causal trees

Causal tree approach extends decision trees to estimate causal effects

Decision trees → Causal trees

- Applying a potential outcome approach to decision trees requires **altering the objective**
  - <u>Objective</u>: Predict, not Y, but the conditional average treatment effect
  - Prefer a tree that **minimizes heterogeneity in leaf-specific treatment effects**

- Yet no "**ground truth**"
  - Estimate an individual treatment effect

Causal tree approach extends decision trees to estimate causal effects

Decision trees → Causal trees

- "**Honest estimation**:" Split the sample into **training data** for generating partitions, and **estimation data** for estimating leaf-specific effects
  - Modified MSE
  - Enables standard asymptotic properties

- Causal trees do not guarantee **unconfoundedness**
  - **Leaf-specific adjustment**, using matching, weighting, or generalized random forests
  - **Sensitivity** analysis

Causal Tree Algorithm Workflow

Input data with selected covariates

Sample splitting for honest estimation

Sample splitting for K-fold cross validation

Resulting partitions from the Causal Tree

Feed estimation sample into terminal leaves

$CATE(\Pi)$ estimation for observations within leaves

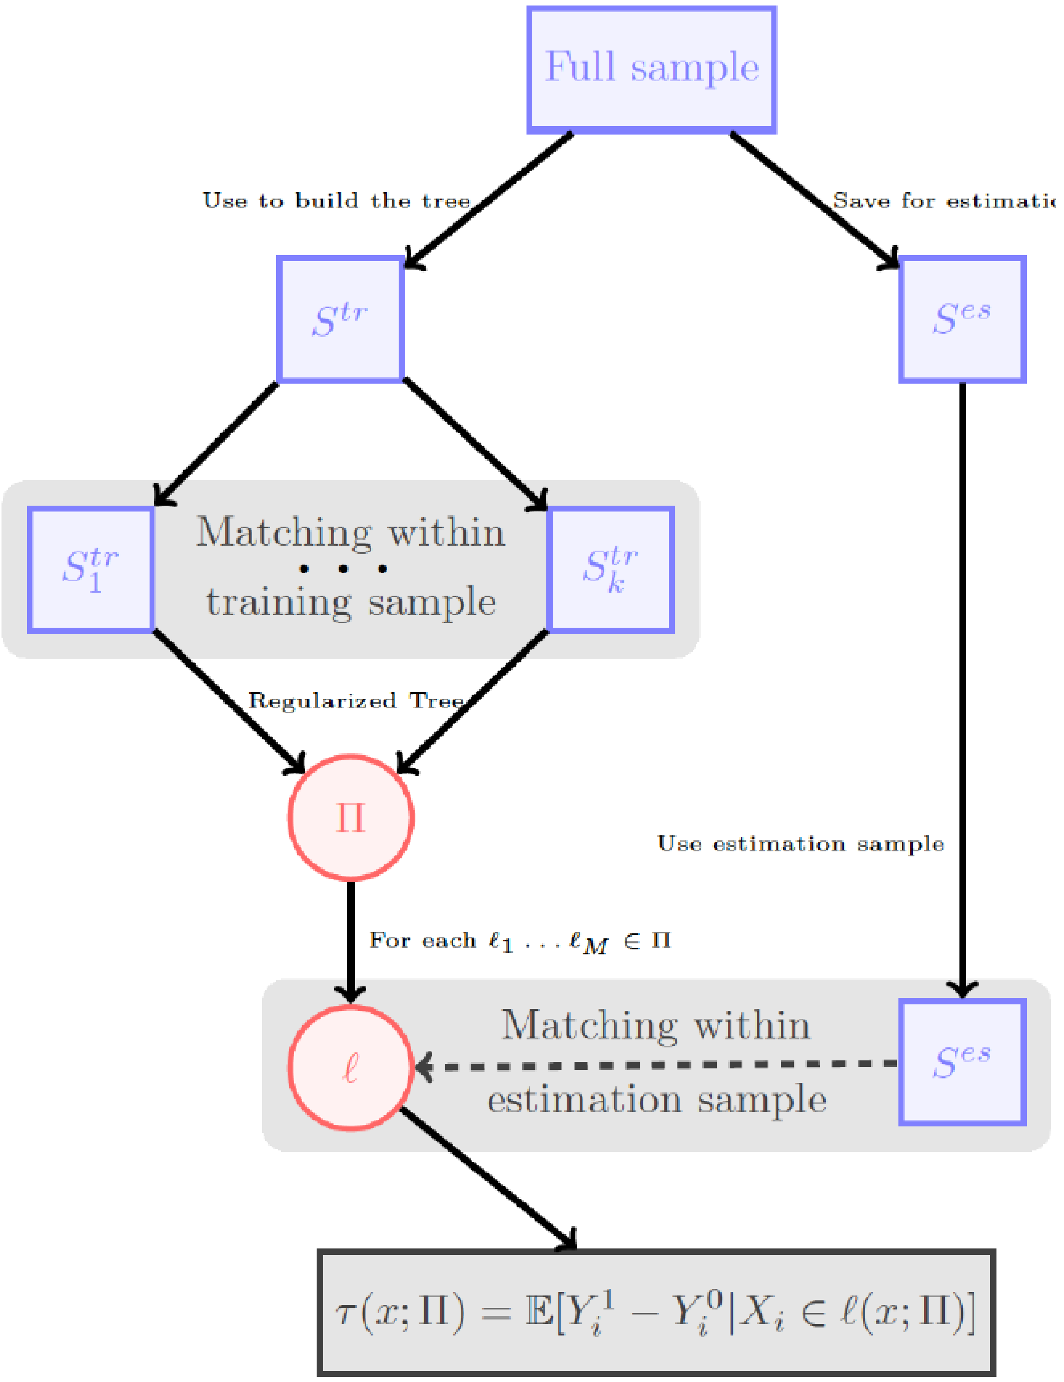

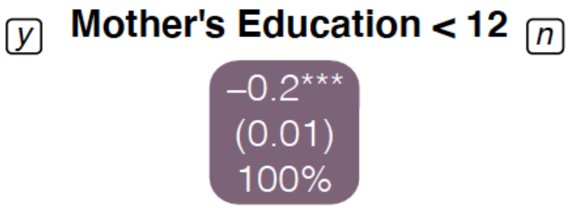
y
Mother's Education < 12
n
−0.2***
(0.01)
100%

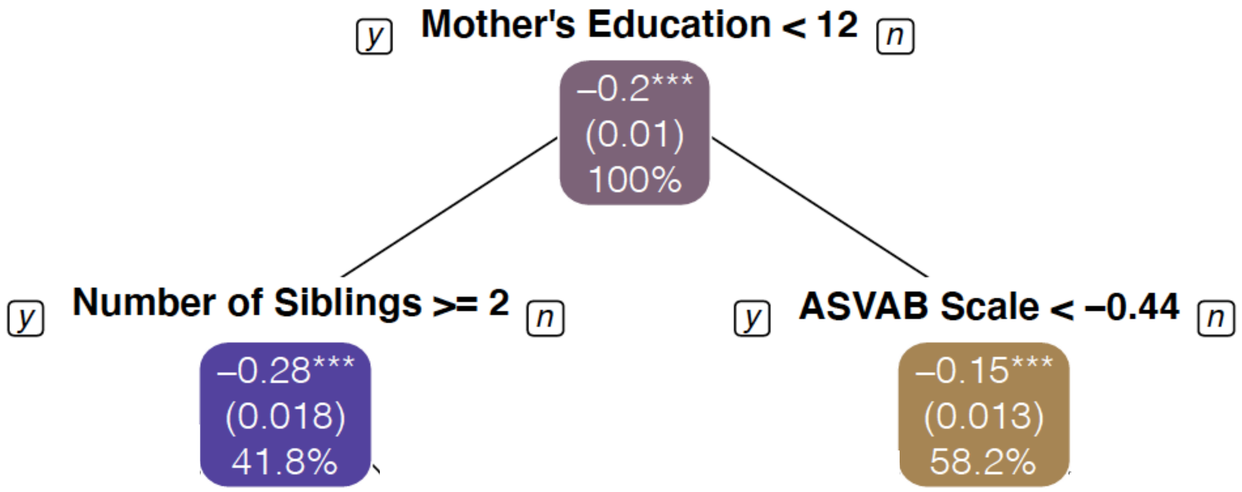

y
Mother's Education < 12
n
–0.2***
(0.01)
100%
y
Number of Siblings >= 2
n
–0.28***
(0.018)
41.8%
y
ASVAB Scale < –0.44
n
–0.15***
(0.013)
58.2%

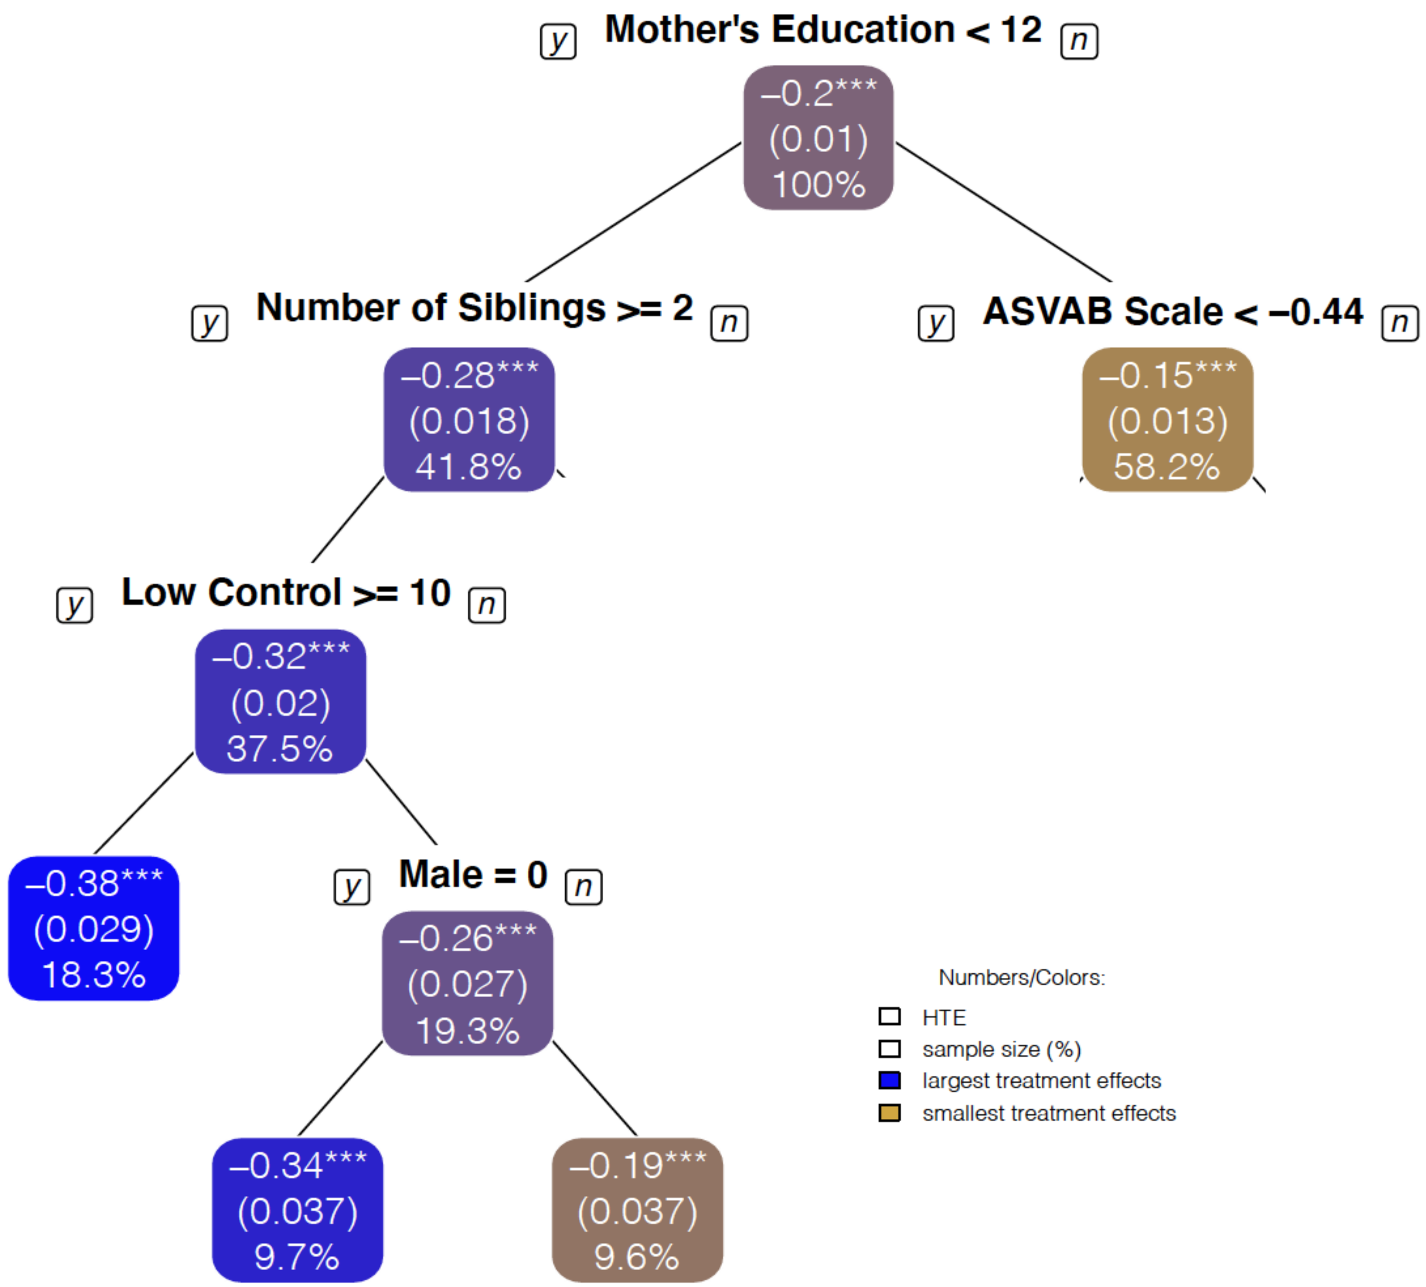

Mother's Education < 12
y
n
–0.2***
(0.01)
100%
Number of Siblings >= 2
–0.28***
(0.018)
41.8%
ASVAB Scale < –0.44
–0.15***
(0.013)
58.2%
Low Control >= 10
–0.32***
(0.02)
37.5%
–0.38***
(0.029)
18.3%
Male = 0
–0.26***
(0.027)
19.3%
–0.34***
(0.037)
9.7%
–0.19***
(0.037)
9.6%
Numbers/Colors:
HTE
sample size (%)
largest treatment effects
smallest treatment effects

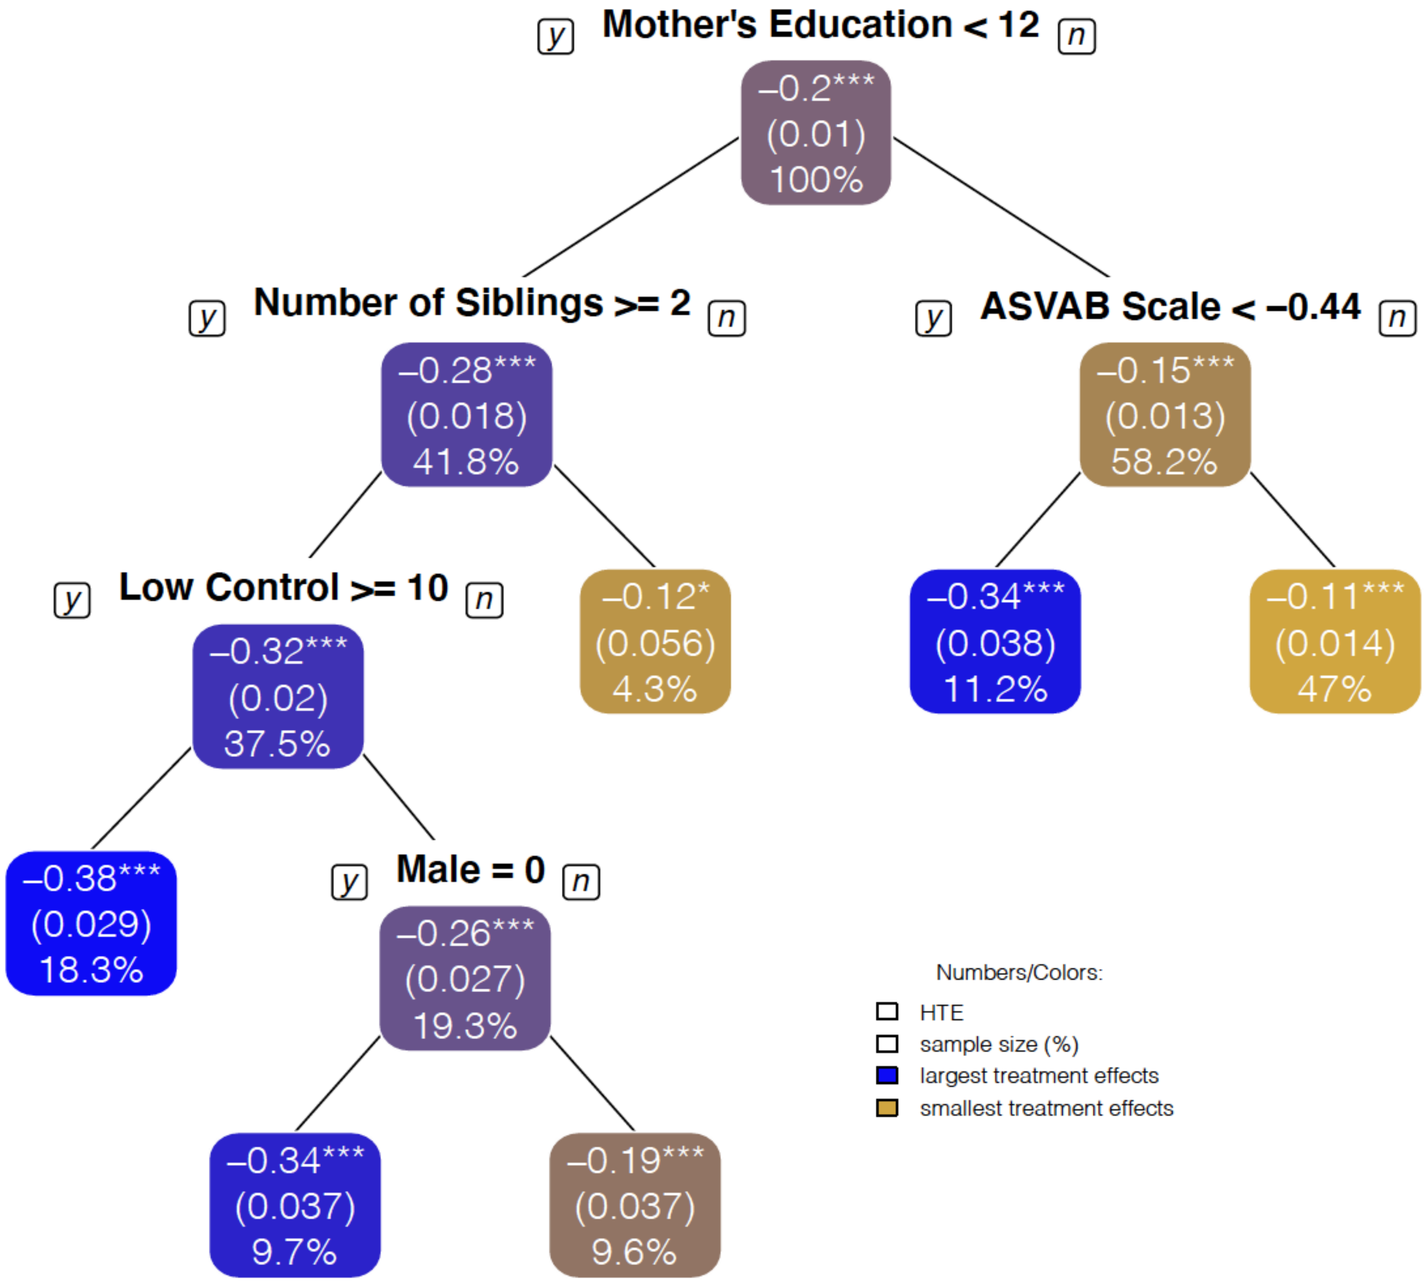

Mother's Education < 12
y
n
–0.2***
(0.01)
100%
Number of Siblings >= 2
–0.28***
(0.018)
41.8%
ASVAB Scale < –0.44
–0.15***
(0.013)
58.2%
Low Control >= 10
–0.32***
(0.02)
37.5%
–0.12*
(0.056)
4.3%
–0.34***
(0.038)
11.2%
–0.11***
(0.014)
47%
–0.38***
(0.029)
18.3%
Male = 0
–0.26***
(0.027)
19.3%
–0.34***
(0.037)
9.7%
–0.19***
(0.037)
9.6%
Numbers/Colors:
HTE
sample size (%)
largest treatment effects
smallest treatment effects

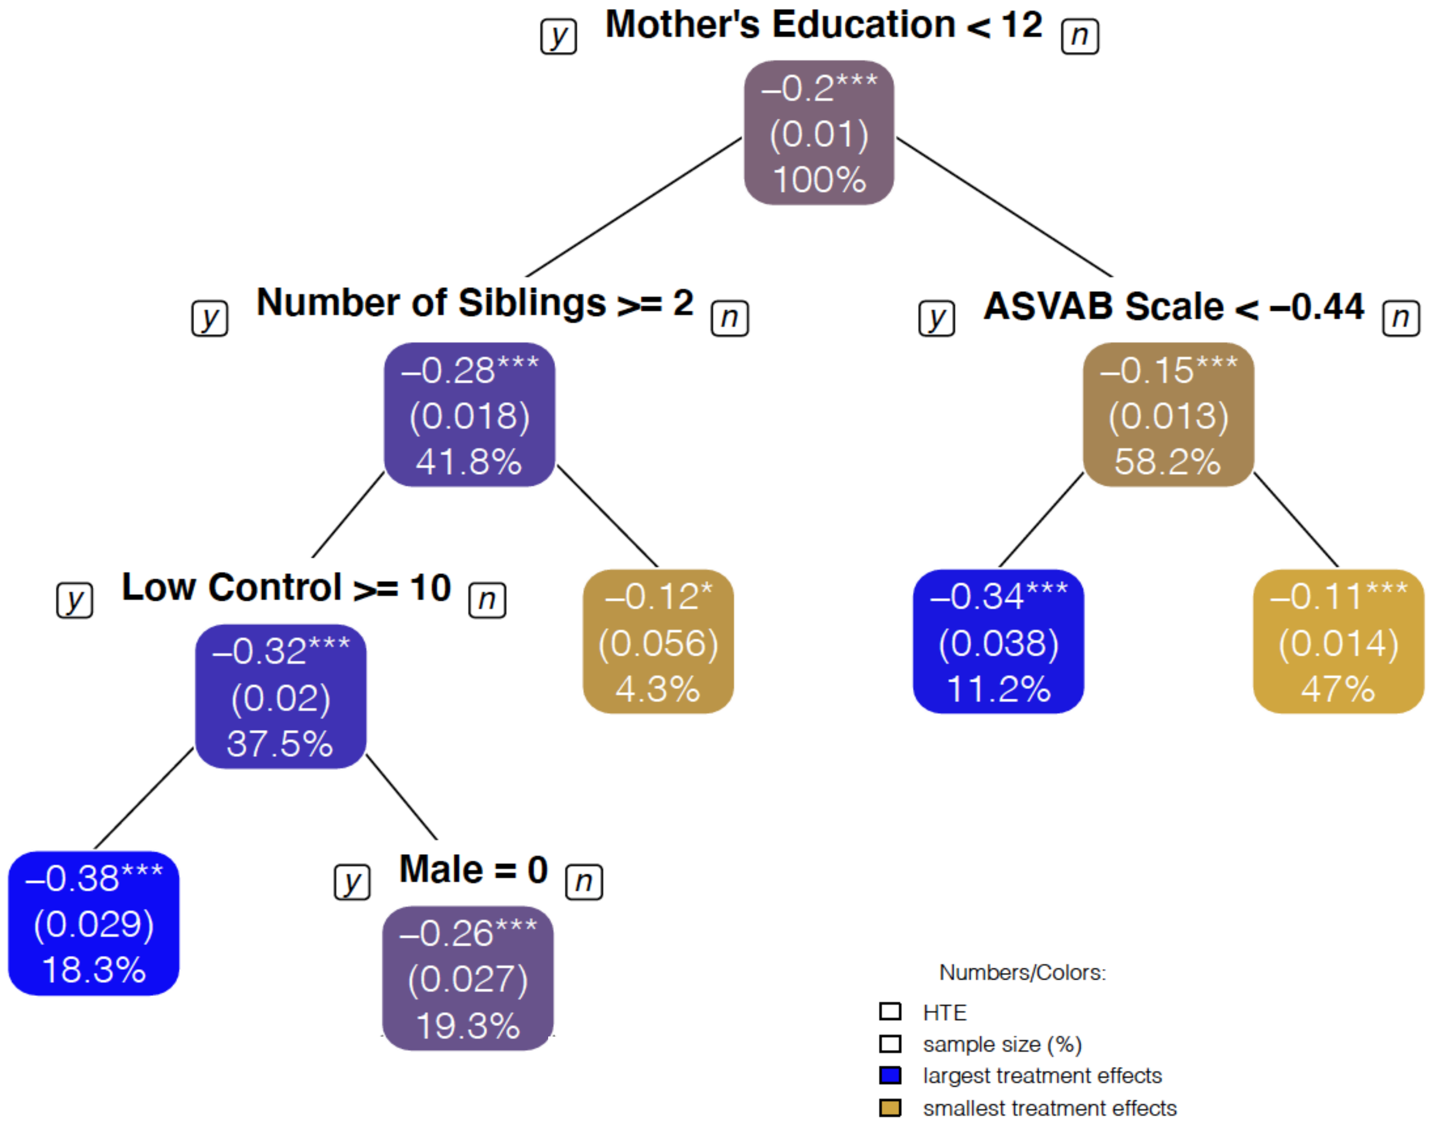

Mother's Education < 12
y
n
–0.2***
(0.01)
100%
Number of Siblings >= 2
–0.28***
(0.018)
41.8%
ASVAB Scale < –0.44
–0.15***
(0.013)
58.2%
Low Control >= 10
–0.32***
(0.02)
37.5%
–0.12*
(0.056)
4.3%
–0.34***
(0.038)
11.2%
–0.11***
(0.014)
47%
–0.38***
(0.029)
18.3%
Male = 0
–0.26***
(0.027)
19.3%
Numbers/Colors:
HTE
sample size (%)
largest treatment effects
smallest treatment effects

Variable Importance Plot Based on Causal Forest

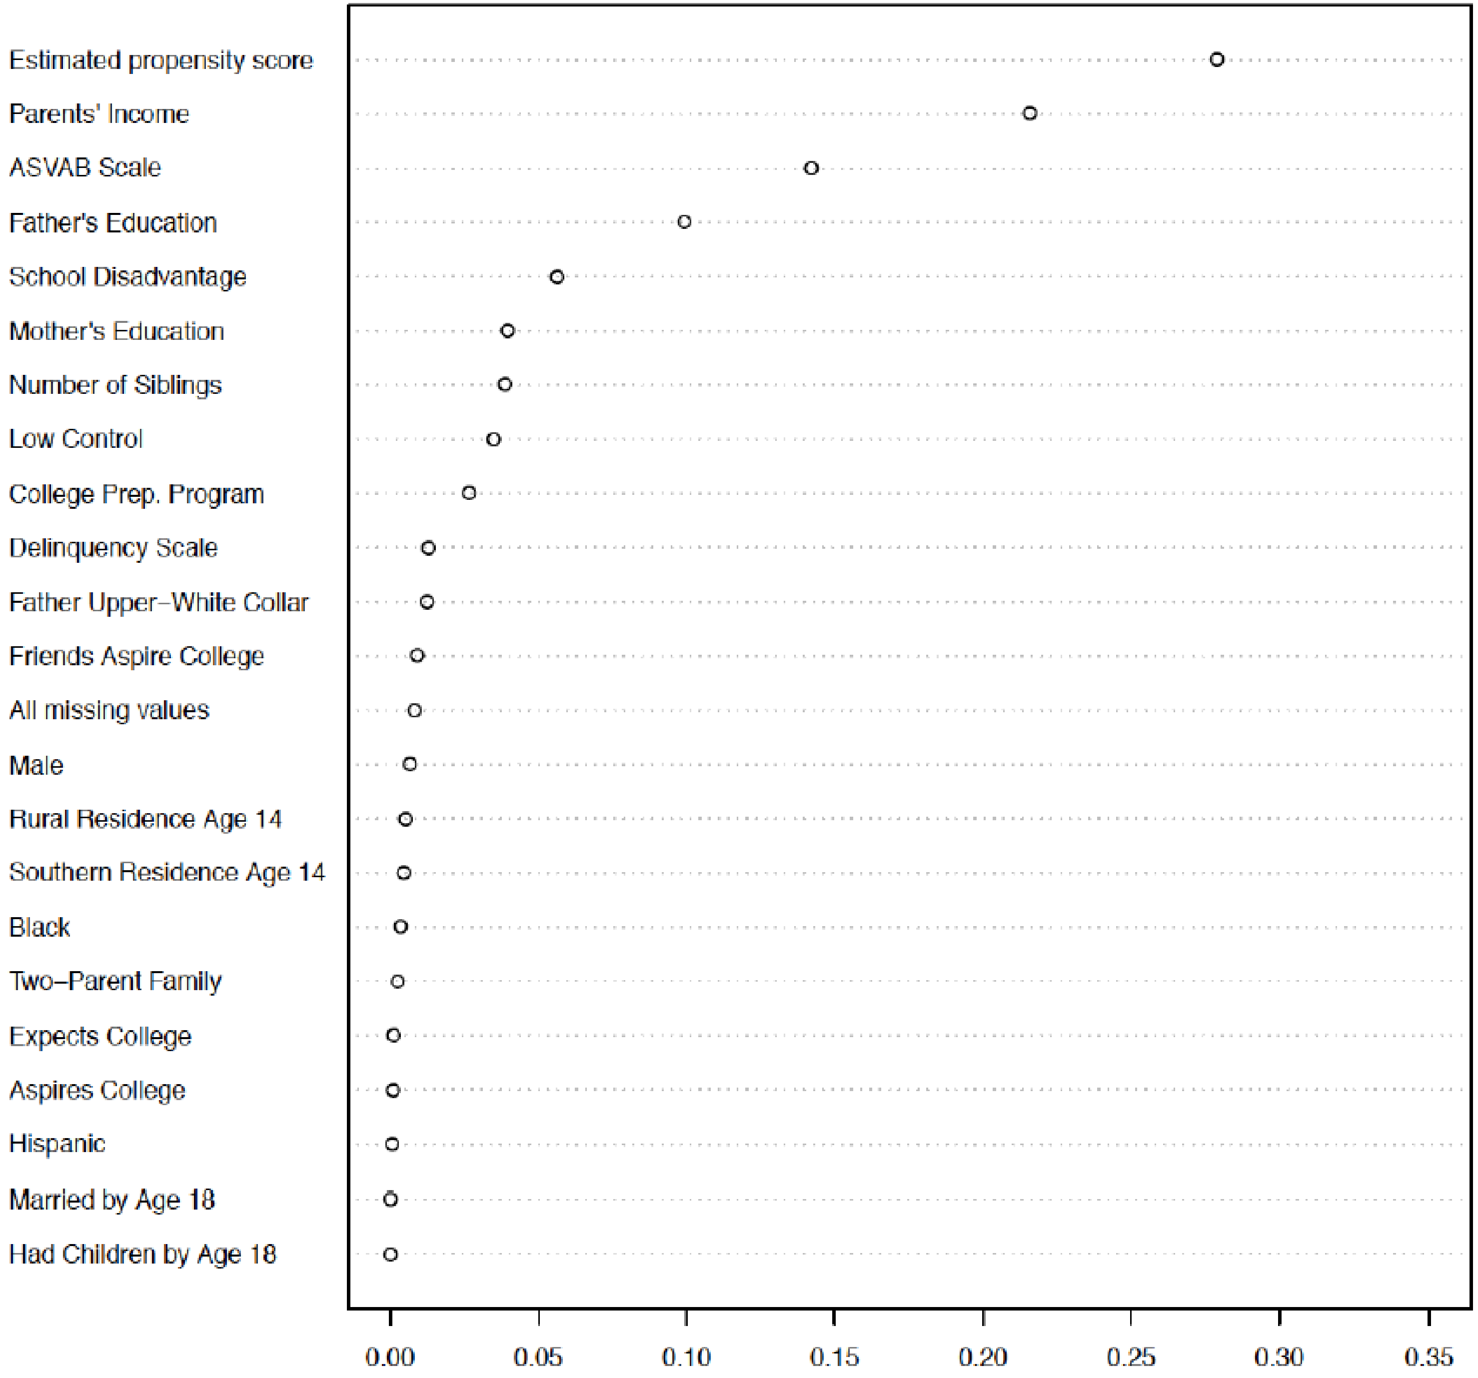

Causal Trees for Uncovering Effect Heterogeneity

- More **disadvantaged** subgroups, or those on the margin of school continuation, have **larger effects** of college on reducing low-wage work
- Groups identified by the **causal tree** are **not identical** to groups identified by **theoretical priors**
- Focus on the **characteristics of the groups identified** by the leaf-specific partitions

Causal Trees for Uncovering Effect Heterogeneity

Leaf 3 – Responsive

Low parental income

High school disadvantage

Low test scores

Low parental education

Majority black or Hispanic

Low social control

Low propensity

Causal Trees for Uncovering Effect Heterogeneity

Leaf 3 – Responsive	Leaf 9 – Responsive
Low parental income	Average income
High school disadvantage	Average school disadvantage
Low test scores	(Very) low test scores
Low parental education	Average parental education
Majority black or Hispanic	Majority white
Low social control	Low social control
Low propensity	Low propensity

Causal Trees for Uncovering Effect Heterogeneity

Leaf 3 – Responsive

Low parental income

High school disadvantage

Low test scores

Low parental education

Majority black or Hispanic

Low social control

Low propensity

Leaf 9 – Responsive

Average income

Average school disadvantage

(Very) low test scores

Average parental education

Majority white

Low social control

Low propensity

Leaf 10 – Least responsive

High income

Low school disadvantage

High test scores

High parental education

Majority white

High social control

High propensity

Causal Inference and Machine Learning: Discussion

- Exciting new developments in causal inference and machine learning
- Uncovering sources of effect heterogeneity is one such development
  - How do we determine meaningful sources of response variation?
    - Causal trees helped highlight particularly responsive subgroups
- The most effective uses of machine learning will likely be in settings where social scientists can define a clear aspect of the problem to outsource to an algorithm

Thank you!

http://www.profjenniebrand.com
brand@soc.ucla.edu