Causal Data Science:
A general framework for data fusion and causal inference

Elias Bareinboim
Columbia University
Twitter: @eliasbareinboim

MIT Graphical Models Workshop
Boston, August, 2019

Causal Data Science:
A general framework for data fusion and causal inference

(On the Causal Foundations of Data Science)

Elias Bareinboim

Columbia University

Twitter: @eliasbareinboim

MIT Graphical Models Workshop

Boston, August, 2019

Causal Data Science:
A general framework for data fusion and causal inference

Elias Bareinboim

Columbia University

Twitter: @eliasbareinboim

MIT Graphical Models Workshop

Boston, August, 2019

The aspirations & challenges of modern Data Science (circa 2019)

The aspirations & challenges of modern Data Science (circa 2019)

- *“The ability to take data — to be able to understand it, to extract value from it, to visualize it, to communicate it - that’s going to be hugely important in the next decades”*

  Hal Varian, chief economist at Google and UC Berkeley Professor of Information Sciences, Business, and Economics.

The aspirations & challenges of modern Data Science (circa 2019)

- *"The ability to take data — to be able to understand it, to extract value from it, to visualize it, to communicate it - that's going to be hugely important in the next decades"*

  Hal Varian, chief economist at Google and UC Berkeley Professor of Information Sciences, Business, and Economics.

- *"Big data is not about the data!"*

  Gary King, Political Scientist, University Professor, Harvard University.

The aspirations & challenges of modern Data Science (circa 2019)

- *"The ability to take data — to be able to understand it, to extract value from it, to visualize it, to communicate it - that's going to be hugely important in the next decades"*

  Hal Varian, chief economist at Google and UC Berkeley Professor of Information Sciences, Business, and Economics.

- *"Big data is not about the data!"*

  Gary King, Political Scientist, University Professor, Harvard University.

- *"Data Science is only as much of a science as it facilitates the interpretation of data - a two body problem, connecting data to reality".*

  Judea Pearl, Professor of Computer Science & Statistics, UCLA.

CURRENT STATE OF AFFAIRS
(REPORT FROM THE TRENCHES)


SEARCH
POWERED BY Google
EDITION: U.S. | INTERNATIONAL | MÉXICO | ARABIC
TV: CNN | CNNI | CNN en Español | HLN
CNN Health
Home
TV & Video
CNN Trends
U.S.
World
Politics
Justice
Entertainment
Tech
Health
Living
Travel
Opinion
iReport
Money
Sports
the chart
Dr. Sanjay Gupta
Children's Health
Expert Doctor Q&A
Sleep
Sex and You
2.3k
Recommend
102
Tweet
31
Share
20
+1
More sharing
Privacy Policy Close
voiceFIVE
Are you aware of any of the following?
UnitedHealthcare
Humana
Aetna
Blue Cross
1 of 5
Only 5 questions? Easy.
Too easy
ADVERTISEMENT
August 15th, 2013
08:00 PM ET
Study: Heavy coffee drinking in people under 55 linked to early death
Dr. Sanjay Gupta


SEARCH
POWERED BY Google
EDITION: U.S.
INTERNATIONAL
MÉXICO
ARABIC
TV: CNN
CNNI
CNN en Español
HLN
CNN Health
Home
TV & Video
CNN Trends
U.S.
World
Politics
Justice
Entertainment
Tech
Health
Living
Travel
Opinion
iReport
Money
Sports
the chart
Dr. Sanjay Gupta
Children's Health
Expert Doctor Q&A
Sleep
Sex and You
2.3k
Recommend
102
Tweet
31
Share
20
+1
More sharing
Privacy Policy Close
voiceFIVE
Are you aware of any of the following?
UnitedHealthcare
Humana
Blue Cross
Aetna
1 of 5
Only 5 questions? Easy.
Too easy
ADVERTISEMENT
Dr. Sanjay Gupta
August 15th, 2013
08:00 PM ET
Study: Heavy coffee drinking in people under 55 linked to early death


EDITION: U.S. | INTERNATIONAL | MÉXICO | ARABIC
TV: CNN | CNNI | CNN en Español | HLN
Home | TV & Video | CNN Trends | U.S. | World
the chart
Dr. Sanjay Gupta
August 15th, 2013
08:00 PM ET
Study: Heavy
linked to early death
Annals of Internal Medicine
Established in 1927 by the American College of Physicians
UCLA Biomedical Lib/Serial
Home | Current Issue | All Issues | Online First | Collections | In the Clinic | Journal Club | CME
17 June 2008, Vol 148, No. 12>
< Previous in this Issue
Next in this Issue >
Email | Share | Get Citation | Figures Slideset (.ppt)
17 June 2008
PDF
The Relationship of Coffee Consumption with Mortality
Rob M. van Dam, PhD; Tricia Y. Li, MD; Fernando Rodriguez-Artalejo, MD, PhD; and
B. Hu, MD, PhD
[+] Article and Author Information
Ann Intern Med. 2008;148(12):904-914. doi:10.7326/0003-4819-148-12-200806170-00003
Article | Figures | Tables | References | Audio/Video | Summary for Patients | Comments (2)
Text Size: A A A
Abstract
Abstract | Context | Contribution | Caution | Methods | Results | Discussion | References
Background: Coffee consumption has been linked to various beneficial and detrimental health effects,
but data on its relation with mortality are sparse.


CBSNEWS
Video
RELAX AND
ENJOY THE LUXURY.
A full-size sedan designed with you in mind.
By MICHELLE CASTILLO / CBS NEWS / February 15, 2013, 3: 36 PM
Alcohol causes 20,000 cancer deaths in the U.S. annually
In Texas it is illegal to take more than three sips of beer at a time while standing. / AP / FILE
41 Comments / Shares / Tweets / Stumble / @ Email
More +
EDITION:
TV: CNN
Home
Dr. Sa
August 15t
08:00 PM E
UCLA Biomedical Lib/Serial
the Clinic
Journal Club
CME
PDF
Mortality
io, MD, PhD; and Frank
Size: A A A
ts (2)
licial and detrimental health effects,


Science News

from universities, journals, and other research organizations

One Drink Of Red Wine Or Alcohol Is Relaxing To Circulation, But Two Drinks Are Stressful

Feb. 13, 2008 — One drink of either red wine or alcohol slightly benefits the heart and blood vessels, but the positive effects on specific biological markers disappear with two drinks, say researchers at the Peter Munk Cardiac Centre of the Toronto General Hospital.

Related Topics

Health & Medicine
- Heart Disease
- Hypertension

Mind & Brain

Articles
- Mediterranean diet
- Drunkenness
- Coronary heart disease

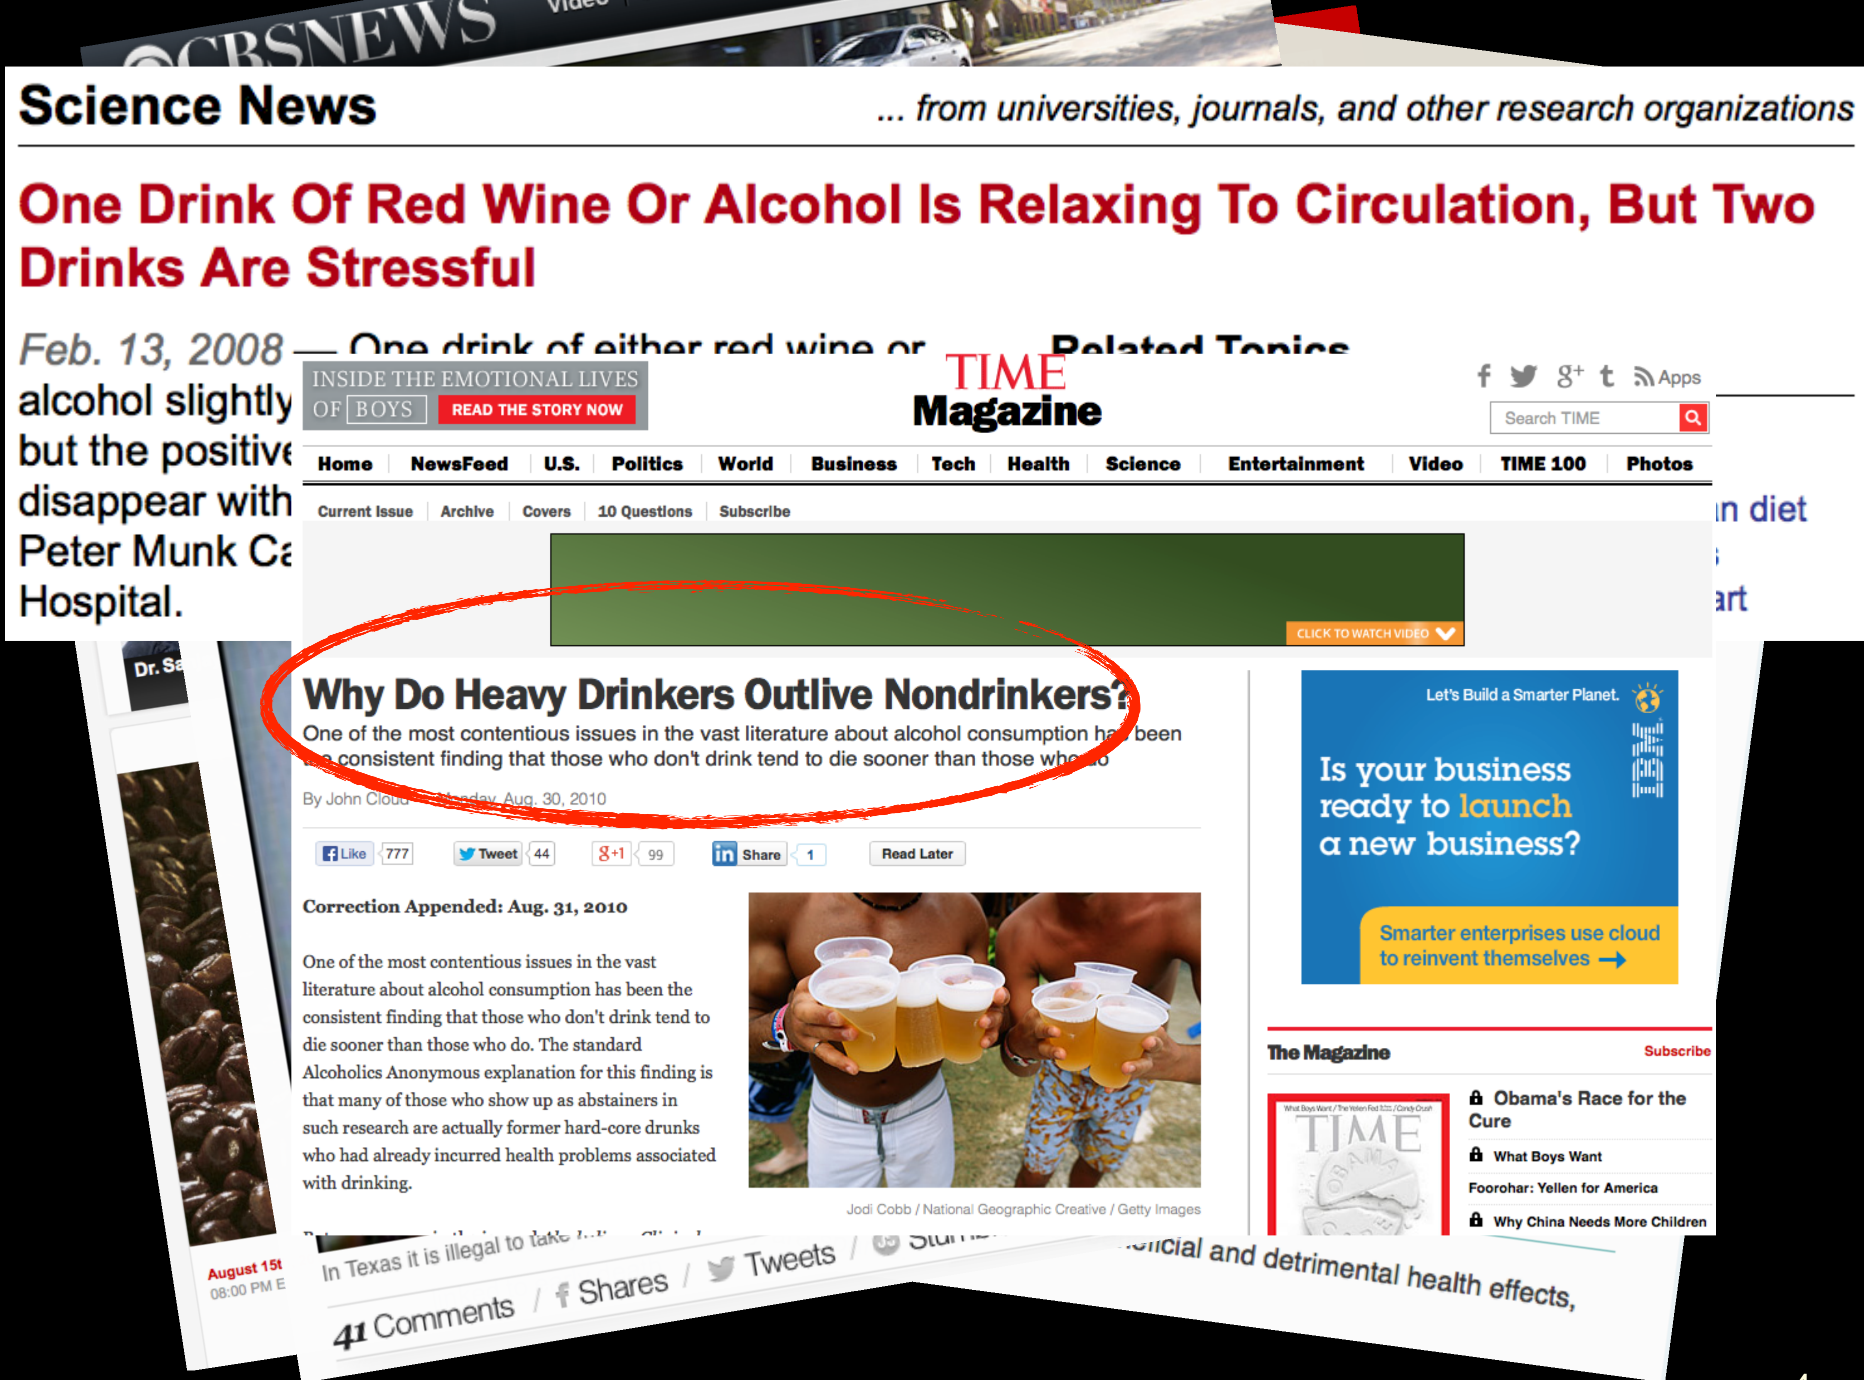

CBSNEWS Video

Science News ... from universities, journals, and other research organizations

One Drink Of Red Wine Or Alcohol Is Relaxing To Circulation, But Two Drinks Are Stressful

Feb. 13, 2008 — One drink of either red wine or alcohol slightly but the positive disappear with Peter Munk Ca Hospital.

Related Topics

INSIDE THE EMOTIONAL LIVES OF BOYS READ THE STORY NOW

TIME Magazine

Search TIME

Home | NewsFeed | U.S. | Politics | World | Business | Tech | Health | Science | Entertainment | Video | TIME 100 | Photos

Current Issue | Archive | Covers | 10 Questions | Subscribe

CLICK TO WATCH VIDEO

Why Do Heavy Drinkers Outlive Nondrinkers?

One of the most contentious issues in the vast literature about alcohol consumption has been the consistent finding that those who don't drink tend to die sooner than those who do

By John Cloud Monday, Aug. 30, 2010

Like 777 | Tweet 44 | g+1 99 | Share 1 | Read Later

Correction Appended: Aug. 31, 2010

One of the most contentious issues in the vast literature about alcohol consumption has been the consistent finding that those who don't drink tend to die sooner than those who do. The standard Alcoholics Anonymous explanation for this finding is that many of those who show up as abstainers in such research are actually former hard-core drunks who had already incurred health problems associated with drinking.

Jodi Cobb / National Geographic Creative / Getty Images

Let's Build a Smarter Planet. IBM

Is your business ready to launch a new business?

Smarter enterprises use cloud to reinvent themselves

The Magazine Subscribe

Obama's Race for the Cure

What Boys Want

Foorohar: Yellen for America

Why China Needs More Children

Dr. Sa

August 15t 08:00 PM E

In Texas it is illegal to take

41 Comments / Shares / Tweets / Stum

...ficial and detrimental health effects,


CBSNEWS
Science News
from universities, journals, and other research organizations
One Drink Of Red Wine Or
To Circulation, But Two
Drinks Are Stressful
Feb. 13, 2008 — One drink of eith
alcohol slightly
but the positive
disappear with
Peter Munk Ca
Hospital.
INSIDE THE EMOTIONAL LIVES OF BOYS
READ THE STORY NOW
Home
NewsFeed
U.S.
Current Issue
Archive
Covers
Why Do Hea
One of the most conter
consistent finding
By John Cloud
Like
777
Correction Appen
One of the most co
literature about a
consistent findi
die sooner tha
Alcoholics An
that many of
such resear
who had already incu
with drinking.
In Texas it is illegal to take
41 Comments
Shares
August 15t
08:00 PM E
The NEW ENGLAND JOURNAL of MEDICINE
HOME
ARTICLES & MULTIMEDIA
ISSUES
SPECIALTIES & TOPICS
FOR AUTHORS
CME
ORIGINAL ARTICLE
Association of Nut Consumption with Total and Cause-Specific Mortality
Ying Bao, M.D., Sc.D., Jiali Han, Ph.D., Frank B. Hu, M.D., Ph.D., Edward L. Giovannucci, M.D., Sc.D., Meir J. Stampfer, M.D., Dr.P.H., Walter C. Willett, M.D., Dr.P.H., and Charles S. Fuchs, M.D., M.P.H.
N Engl J Med 2013; 369:2001-2011 | November 21, 2013 | DOI: 10.1056/NEJMoa1307352
Abstract
Article
References
Share:
BACKGROUND
Increased nut consumption has been associated with a reduced risk of major chronic diseases, including cardiovascular disease and type 2 diabetes mellitus. However, the association between nut consumption and mortality remains unclear.
Full Text of Background...
METHODS
We examined the association between nut consumption and
MEDIA IN THIS ARTICLE
Video
Nuts

WHAT’S GOING ON HERE?

WHAT’S GOING ON HERE? 🤷‍♂️

WHAT’S GOING ON HERE?

Eli’s thesis: Mismatch between the type of data collected & desired claim.

TASKS

TASKS

- Develop machinery (language, conditions, and algorithms) for performing two tasks:

TASKS

- Develop machinery (language, conditions, and algorithms) for performing two tasks:

  1. Learning about population-level interventions by cohesively combining multiple heterogenous datasets (NeurIPS'14, PNAS'16, AAAI'17, UAI'18).

TA

- Develop machinery
  and algorithms) for

  1. Learning about p
     interventions by cohesively combining
     multiple heterogenous datasets
     (NeurIPS'14, PNAS'16, AAAI'17, UAI'18).

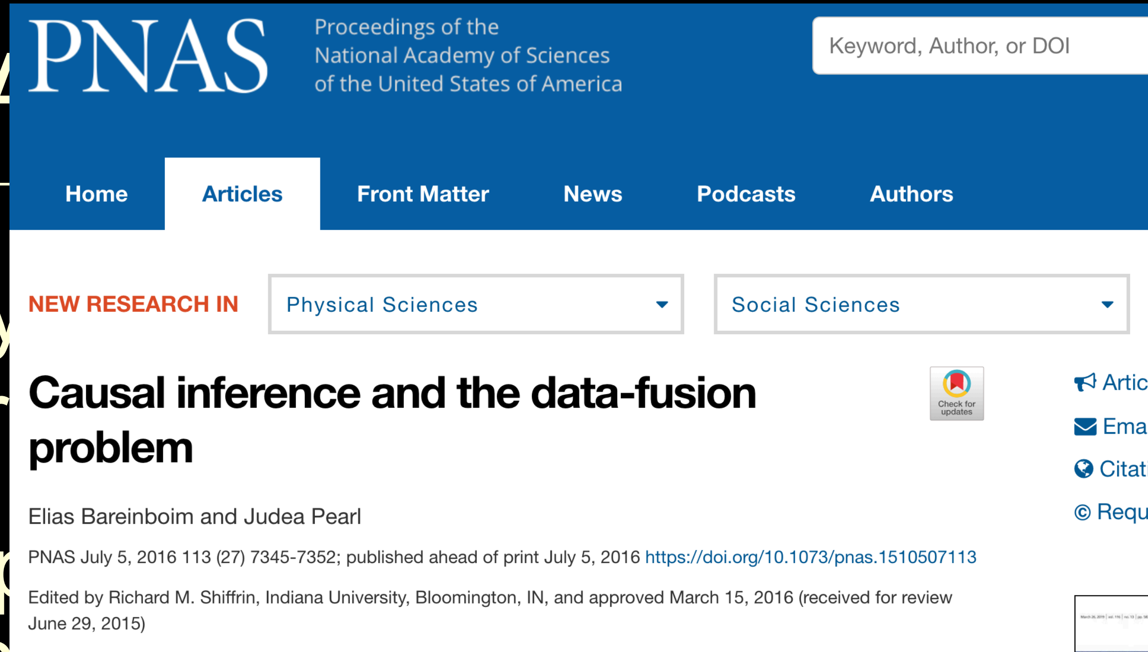

PNAS Proceedings of the National Academy of Sciences of the United States of America

Keyword, Author, or DOI

Home | Articles | Front Matter | News | Podcasts | Authors

NEW RESEARCH IN Physical Sciences | Social Sciences

Causal inference and the data-fusion problem

Elias Bareinboim and Judea Pearl

PNAS July 5, 2016 113 (27) 7345-7352; published ahead of print July 5, 2016 https://doi.org/10.1073/pnas.1510507113

Edited by Richard M. Shiffrin, Indiana University, Bloomington, IN, and approved March 15, 2016 (received for review June 29, 2015)

TASKS

- Develop machinery (language, conditions, and algorithms) for performing two tasks:

  1. Learning about population-level interventions by cohesively combining multiple heterogenous datasets

     (NeurIPS'14, PNAS'16, AAAI'17, UAI'18).

TASKS

- Develop machinery (language, conditions, and algorithms) for performing two tasks:

1. Learning about population-level interventions by cohesively combining multiple heterogenous datasets (NeurIPS'14, PNAS'16, AAAI'17, UAI'18).

2. Deciding individual-level treatments by leveraging population-level knowledge (NeurIPS'15, ICML'17, NeurIPS'18, AAAI'19).

DATA, DATA, DATA…
Challenge: “All data is not created equal”

DATA, DATA, DATA…
Challenge: “All data is not created equal”

- Key observation. There’s a lot of data out there, but this data is almost invariably collected…
  - under different experimental conditions,

DATA, DATA, DATA…
Challenge: "All data is not created equal"

- Key observation. There's a lot of data out there, but this data is almost invariably collected…

  - under different experimental conditions,
  - the underlying populations are different,

DATA, DATA, DATA…
Challenge: “All data is not created equal”

- Key observation. There’s a lot of data out there, but this data is almost invariably collected…
  - under different experimental conditions,
  - the underlying populations are different,
  - the sampling procedure is not random,

DATA, DATA, DATA…
Challenge: "All data is not created equal"

- Key observation. There's a lot of data out there, but this data is almost invariably collected…
  - under different experimental conditions,
  - the underlying populations are different,
  - the sampling procedure is not random,
  - the treatment assignment is not random,

DATA, DATA, DATA…
Challenge: “All data is not created equal”

- Key observation. There’s a lot of data out there, but this data is almost invariably collected…
  - under different experimental conditions,
  - the underlying populations are different,
  - the sampling procedure is not random,
  - the treatment assignment is not random,
  - many variables are not measured.

DATA, DATA, DATA…
Challenge: "All data is not created equal"

- Key observation. There's a lot of data out there, but this data is almost invariably collected…

  - under different experimental conditions,
  - the underlying populations are different,
  - the sampling procedure is not random,
  - the treatment assignment is not random,
  - many variables are not measured.

- In words, the collected data is messy, and rarely perfectly matches the inferential target.

DATA, DATA, DATA…
Challenge: "All data is not created equal"

- Key observation. There's a lot of data out there, but this data is almost invariably collected…

  - under different experimental conditions,
  - the underlying populations are different,
  - the sampling procedure is not random,
  - the treatment assignment is not random,
  - many variables are not measured.

- In words, the collected data is messy, and rarely perfectly matches the inferential target.
- Positive: All these dimensions are now formalized.

DATA, DATA, DATA…
Challenge: “All data is not created equal”

- Key observation. There’s a lot of data out there, but this data is almost invariably collected…

  - under different experimental conditions,
  - the underlying populations are different,
  - the sampling procedure is not random,
  - the treatment assignment is not random,
  - many variables are not measured.

- In words, the collected data is messy, and rarely perfectly matches the inferential target.
- Positive: All these dimensions are now formalized.
- And there are conditions and algorithms to decide what is “entailed” from a certain data collection.

HETEROGENOUS DATASETS

Target

$Q = P^*(y \mid do(x))$

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$
vs $P^*(y \mid x)$

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$

HETEROGENOUS DATASETS

Dataset 1

Dataset 2

…

Dataset n

Target

$Q = P^*(y \mid do(x))$

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$

Dataset 1 | Dataset 2 | … | Dataset n

	Dataset 1	Dataset 2	Dataset n
Population	Los Angeles	New York	Texas
Obs. / Exp.	Experimental	Observational	Experimental
Treat. Assign.	Randomized Z_1	-	Randomized Z_2
Sampling	Selection on Age	Selection on SES	-
Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$

Dataset 1 | Dataset 2 | … | Dataset n

d_1	Population	Los Angeles	New York	Texas
	Obs. / Exp.	Experimental	Observational	Experimental
	Treat. Assign.	Randomized Z_1	-	Randomized Z_2
	Sampling	Selection on Age	Selection on SES	-
	Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

HETEROGENOUS DATASETS

Target

$Q = P^*(y \mid do(x))$

Dataset 1 | Dataset 2 | … | Dataset n

d_1	Population	Los Angeles	New York	Texas
	Obs. / Exp.	Experimental	Observational	Experimental
	Treat. Assign.	Randomized Z_1	-	Randomized Z_2
	Sampling	Selection on Age	Selection on SES	-
	Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

HETEROGENOUS DATASETS

Target

$Q = P^*(y \mid do(x))$

Dataset 1 Dataset 2 … Dataset n

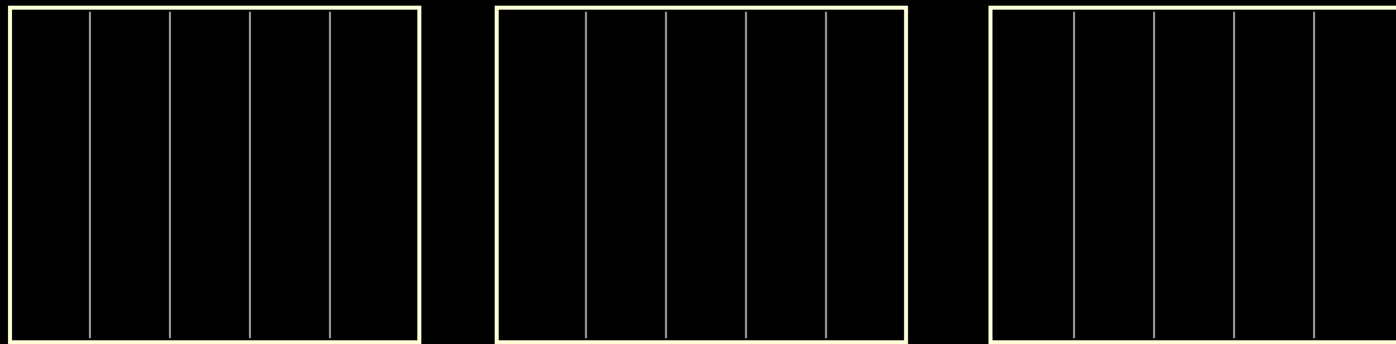

d_1	Population	Los Angeles	New York	Texas
d_2	Obs. / Exp.	Experimental	Observational	Experimental
	Treat. Assign.	Randomized Z_1	-	Randomized Z_2
	Sampling	Selection on Age	Selection on SES	-
	Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$

Dataset 1   Dataset 2   …   Dataset n

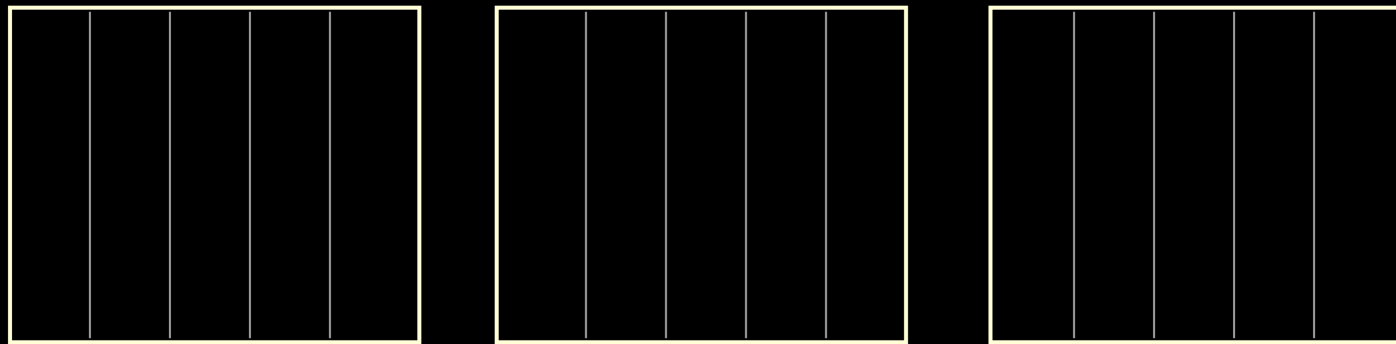

d_1	Population	Los Angeles	New York	Texas
d_2	Obs. / Exp.	Experimental	Observational	Experimental
	Treat. Assign.	Randomized Z_1	-	Randomized Z_2
	Sampling	Selection on Age	Selection on SES	-
	Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$

Dataset 1 | Dataset 2 | … | Dataset n

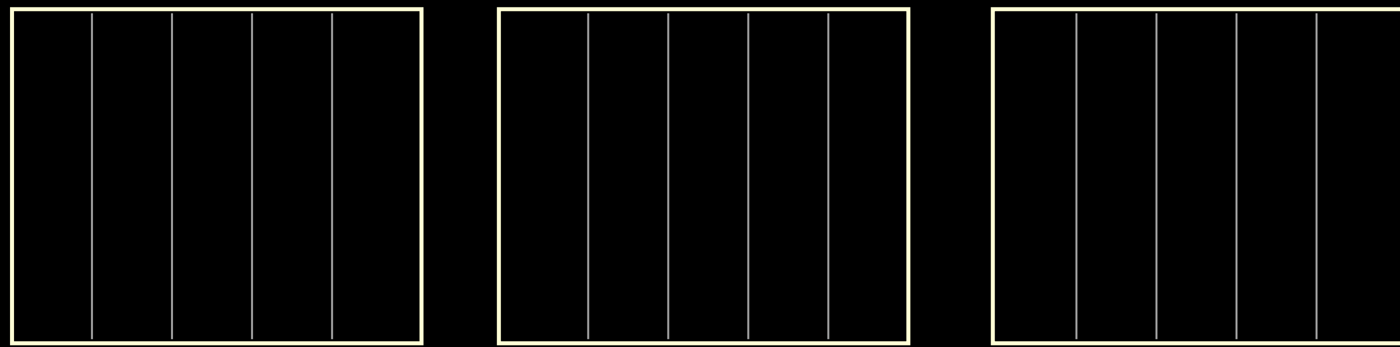

d_1	Population	Los Angeles	New York	Texas
d_2	Obs. / Exp.	Experimental	Observational	Experimental
	Treat. Assign.	Randomized Z_1	-	Randomized Z_2
d_3	Sampling	Selection on Age	Selection on SES	-
	Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$

Dataset 1 | Dataset 2 | … | Dataset n

d_1	Population	Los Angeles	New York	Texas
d_2	Obs. / Exp.	Experimental	Observational	Experimental
	Treat. Assign.	Randomized Z_1	-	Randomized Z_2
d_3	Sampling	Selection on Age	Selection on SES	-
	Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$

Dataset 1 | Dataset 2 | … | Dataset n

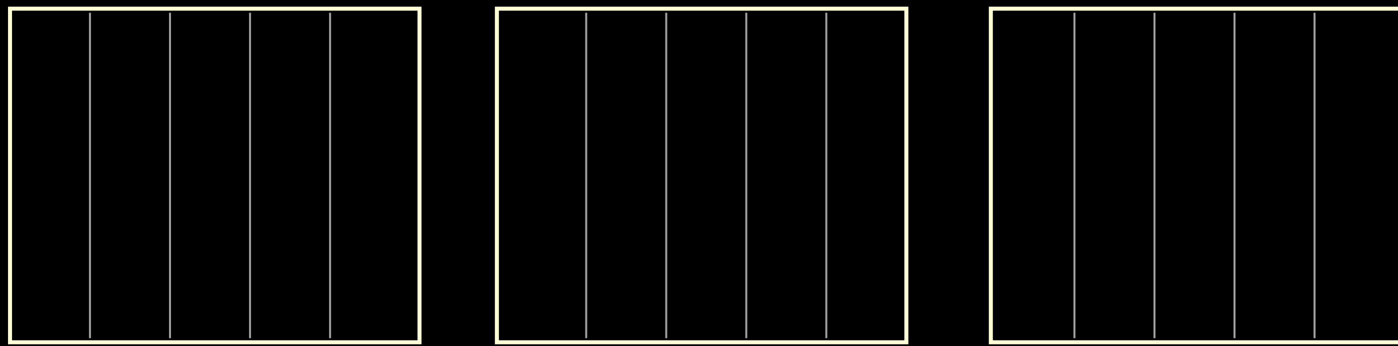

d_1	Population	Los Angeles	New York	Texas
d_2	Obs. / Exp.	Experimental	Observational	Experimental
	Treat. Assign.	Randomized Z_1	-	Randomized Z_2
d_3	Sampling	Selection on Age	Selection on SES	-
d_4	Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

HETEROGENOUS DATASETS

Target
$Q = P^*(y \mid do(x))$

Dataset 1 Dataset 2 … Dataset n

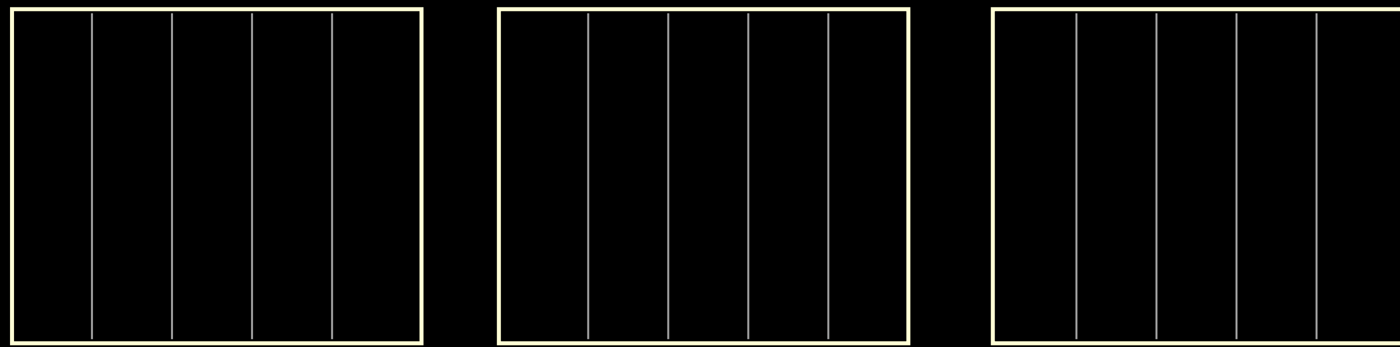

d_1	Population	Los Angeles	New York	Texas
d_2	Obs. / Exp.	Experimental	Observational	Experimental
	Treat. Assign.	Randomized Z_1	-	Randomized Z_2
d_3	Sampling	Selection on Age	Selection on SES	-
d_4	Measured	X_1, Z_1, W, M, Y_1	X_1, X_2, Z_1, N, Y_2	X_2, Z_1, W, L, M, Y_1

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

a. Statistics − Descriptive

Samples(Obs) → Distrib(Obs)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

a. Statistics − Descriptive

Bernoulli, Poisson, Kolmogorov

Samples(Obs) $\rightarrow$ Distrib(Obs)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

a. Statistics – Descriptive

Bernoulli, Poisson, Kolmogorov

Samples(Obs) → Distrib(Obs)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

a. Statistics − Descriptive

Bernoulli, Poisson, Kolmogorov

Samples(Obs) → Distrib(Obs)

b. Statistics − Experimental

Samples(do(X)) → Distrib(do(X))

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

a. Statistics − Descriptive Bernoulli, Poisson, Kolmogorov

Samples(Obs) → Distrib(Obs)

b. Statistics − Experimental Fisher, Cox, Cochran

Samples(do(X)) → Distrib(do(X))

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

a. Statistics − Descriptive

Bernoulli, Poisson, Kolmogorov

Samples(Obs) → Distrib(Obs)

b. Statistics − Experimental

Fisher, Cox, Cochran

Samples(do(X)) → Distrib(do(X))

c. *Causal Inference from Observational Studies*

Distrib(Obs) → Distrib(do(X))

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

a. Statistics − Descriptive

Bernoulli, Poisson, Kolmogorov

Samples(Obs) → Distrib(Obs)

b. Statistics − Experimental

Fisher, Cox, Cochran

Samples(do(X)) → Distrib(do(X))

c. *Causal Inference from Observational Studies*

Distrib(Obs) → Distrib(do(X))

"Causal Inference" from ~1974-2011 was mainly concerned with confounding bias, and has been "solved" by Rubin, Robins, Dawid, Pearl.

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Elias, are you saying we are done?

Some common inferences in scientific circles, AI, and machine learning involving some standard assumptions:

Tasks

a. Statistics − Descriptive Bernoulli, Poisson, Kolmogorov

Samples(Obs) → Distrib(Obs)

b. Statistics − Experimental Fisher, Cox, Cochran

Samples(do(X)) → Distrib(do(X))

c. *Causal Inference from Observational Studies*

Distrib(Obs) → Distrib(do(X))

"Causal Inference" from ~1974-2011 was mainly concerned with confounding bias, and has been "solved" by Rubin, Robins, Dawid, Pearl.

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Elias, are you saying we are done?

Some commo...

machine ... ons:

Also, this is “old stuff”, where does causal data science comes into play?

Tasks

a. S...

Samples(Obs) → Distrib(Obs)

b. Statistics – Experimental

Fisher, Cox, Cochran

Samples(do(X)) → Distrib(do(X))

c. *Causal Inference from Observational Studies*

Distrib(Obs) → Distrib(do(X))

"Causal Inference” from ~1974-2011 was mainly concerned with confounding bias, and has been “solved” by Rubin, Robins, Dawid, Pearl.

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

a. Statistics − Descriptive
Samples(Obs) → Distrib(Obs)

b. Statistics − Experimental
Samples(do(X)) → Distrib(do(X))

c. *Causal Inference from Observational Studies*
Distrib(Obs) → Distrib(do(X))

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

a. Statistics − Descriptive
Samples(Obs) → Distrib(Obs)

b. Statistics − Experimental
Samples(do(X)) → Distrib(do(X))

c. *Causal Inference from Observational Studies*
Distrib(Obs) → Distrib(do(X))

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

a. Statistics − Descriptive
Samples(Obs) → Distrib(Obs)

b. Statistics − Experimental
Samples(do(X)) → Distrib(do(X))

c. *Causal Inference from Observational Studies*
Distrib(Obs) → Distrib(do(X))

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

a. Statistics - Descriptive
(d_1, Samples(Obs), d_3, d_4) → (d_1, Distrib(Obs), d_3, d_4)

b. Statistics - Experimental
(d_1, Samples(do(X)), d_3, d_4) → (d_1, Distrib(do(X)), d_3, d_4)

c. *Causal Inference from Observational Studies*
(d_1, Obs, d_3, d_4) → (d_1, do(X), d_3, d_4)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

a. Statistics - Descriptive
(d_1, Samples(Obs), d_3, d_4) → (d_1, Distrib(Obs), d_3, d_4)

b. Statistics - Experimental
(d_1, Samples(do(X)), d_3, d_4) → (d_1, Distrib(do(X)), d_3, d_4)

c. *Causal Inference from Observational Studies*
(d_1, Obs, d_3, d_4) → (d_1, do(X), d_3, d_4)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

a. Statistics - Descriptive
$(d_1, \text{Samples}(\text{Obs}), d_3, d_4) \rightarrow (d_1, \text{Distrib}(\text{Obs}), d_3, d_4)$

b. Statistics - Experimental
$(d_1, \text{Samples}(do(X)), d_3, d_4) \rightarrow (d_1, \text{Distrib}(do(X)), d_3, d_4)$

c. *Causal Inference from Observational Studies*
$(d_1, \text{Obs}, d_3, d_4) \rightarrow (d_1, do(X), d_3, d_4)$

Do these dimensions exhaust all possible data collection modes?

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

1. *Causal Inference from Obs Studies*
   $(d_1, \text{Obs}, d_3, d_4) \rightarrow (d_1, do(X), d_3, d_4)$

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

1. *Causal Inference from Obs Studies* Rubin, Robins, Dawid, Pearl

$(d_1, \text{Obs}, d_3, d_4) \rightarrow (d_1, do(X), d_3, d_4)$

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

1. *Causal Inference from Obs Studies* Rubin, Robins, Dawid, Pearl
   (d_1, Obs, d_3, d_4) → (d_1, do(X), d_3, d_4)

2. *Experimental Inference* (generalized IVs)
   (d_1, do(Z), d_3, d_4) → (d_1, do(X), d_3, d_4)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

1. *Causal Inference from Obs Studies* Rubin, Robins, Dawid, Pearl
$(d_1, \text{Obs}, d_3, d_4) \rightarrow (d_1, do(X), d_3, d_4)$

2. *Experimental Inference* (generalized IVs) P. & S. Wright, 1928
$(d_1, do(Z), d_3, d_4) \rightarrow (d_1, do(X), d_3, d_4)$

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

1. *Causal Inference from Obs Studies* Rubin, Robins, Dawid, Pearl
   $(d_1, \text{Obs}, d_3, d_4) \rightarrow (d_1, do(X), d_3, d_4)$

2. *Experimental Inference* (generalized IVs) P. & S. Wright, 1928
   $(d_1, do(Z), d_3, d_4) \rightarrow (d_1, do(X), d_3, d_4)$

3. *Sampling Selection Bias*
   $(d_1, d_2, \text{Select}(Age), d_4) \rightarrow (d_1, d_2, \{\}, d_4)$

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

1. *Causal Inference from Obs Studies* Rubin, Robins, Dawid, Pearl
   (d_1, Obs, d_3, d_4) ⟶ (d_1, do(X), d_3, d_4)

2. *Experimental Inference* (generalized IVs) P. & S. Wright, 1928
   (d_1, do(Z), d_3, d_4) ⟶ (d_1, do(X), d_3, d_4)

3. *Sampling Selection Bias* Heckman, 1976
   (d_1, d_2, Select(Age), d_4) ⟶ (d_1, d_2, {}, d_4)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

1. *Causal Inference from Obs Studies* Rubin, Robins, Dawid, Pearl
   $(d_1, \text{Obs}, d_3, d_4) \longrightarrow (d_1, do(X), d_3, d_4)$

2. *Experimental Inference* (generalized IVs) P. & S. Wright, 1928
   $(d_1, do(Z), d_3, d_4) \longrightarrow (d_1, do(X), d_3, d_4)$

3. *Sampling Selection Bias* Heckman, 1976
   $(d_1, d_2, \text{Select}(Age), d_4) \longrightarrow (d_1, d_2, \{\}, d_4)$

4. *Transportability* (External Validity)
   $(\text{Bonobos}, d_2, d_3, d_4) \longrightarrow (\text{Humans}, d_2, d_3, d_4)$

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Tasks

1. *Causal Inference from Obs Studies* Rubin, Robins, Dawid, Pearl
   $(d_1, \text{Obs}, d_3, d_4) \longrightarrow (d_1, do(X), d_3, d_4)$
2. *Experimental Inference* (generalized IVs) P. & S. Wright, 1928
   $(d_1, do(Z), d_3, d_4) \longrightarrow (d_1, do(X), d_3, d_4)$
3. *Sampling Selection Bias* Heckman, 1976
   $(d_1, d_2, \text{Select}(\text{Age}), d_4) \longrightarrow (d_1, d_2, \{\}, d_4)$
4. *Transportability* (External Validity) Shadish, Cook, Campbell, 1976
   $(\text{Bonobos}, d_2, d_3, d_4) \longrightarrow (\text{Humans}, d_2, d_3, d_4)$

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Eli: Special cases of these dimensions have been addressed in the literature, mostly in isolation, and under very special parametric conditions.

In practice, they appear together in what I like to call **causal data science**,

$$\{ (d_1, d_2, d_3, d_4) \} \rightarrow (d^*_1, d^*_2, d^*_3, d^*_4)$$

Tasks

4. *Transportability* (External Validity) Shadish, Cook, Campbell, 1976

(Bonobos, d_2, d_3, d_4) $\rightarrow$ (Humans, d_2, d_3, d_4)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)

Eli: Special cases of these dimensions have been addressed in the literature, mostly in isolation, and under very special parametric conditions.

In practice, they appear together in what I like to call **causal data science**,

$$\{ (d_1, d_2, d_3, d_4) \} \rightarrow (d^*_1, d^*_2, d^*_3, d^*_4)$$

Tasks

4. *Transportability* (External Validity) Shadish, Cook, Campbell, 1976

(Bonobos, d_2, d_3, d_4) $\rightarrow$ (Humans, d_2, d_3, d_4)

CAUSAL DATA SCIENCE - DIMENSIONS & TASKS

Description of each data collection: tuple (d_1, d_2, d_3, d_4)
(population, obs./exp., sampling, measure.)

Eli: Special cases of these dimensions have been addressed in the literature, mostly in isolation, and under very special parametric conditions.

In practice, they appear together in what I like to call **causal data science**,

$$\{ (d_1, d_2, d_3, d_4) \} \rightarrow (d^*_1, d^*_2, d^*_3, d^*_4)$$

Tasks

4. *Transportability* (External Validity) Shadish, Cook, Campbell, 1976

(Bonobos, d_2, d_3, d_4) $\rightarrow$ (Humans, d_2, d_3, d_4)

TRANSPORTABILITY:

EXTRAPOLATION & ROBUSTNESS OF CAUSAL CLAIMS

TRANSPORTABILITY - PROBLEM STATEMENT

Question:

Is it possible to compute the effect of X on Y in a target environment $\prod^*$ (where no experiments are feasible), using experimental findings from a different environment $\prod$?

Answer: Sometimes yes.

TRANSPORTABILITY - PROBLEM STATEMENT

(or, external validity, robustness, generalizability)

Question:

Is it possible to compute the effect of X on Y in a target environment $\prod^*$ (where no experiments are feasible), using experimental findings from a different environment $\prod$?

Answer: Sometimes yes.

TRANSPORTABILITY - PROBLEM STATEMENT

(or, external validity, robustness, generalizability)

Question:

Is it possible to compute the effect of X on Y in a target environment $\prod^*$ (where no experiments are feasible), using experimental findings from a different environment $\prod$?

Answer: Sometimes yes.

Our goal is to formally characterize when and how.

MOVING FROM THE “LAB” TO THE “REAL WORLD”...

Lab

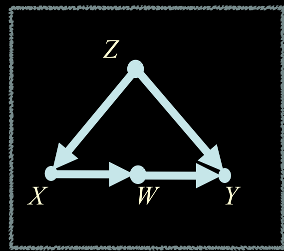

MOVING FROM THE “LAB” TO THE “REAL WORLD”...

Lab

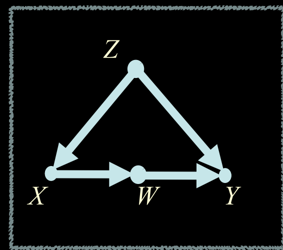

* The lab stands for any environment, population, domain, setting.

MOVING FROM THE "LAB" TO THE "REAL WORLD"...

Lab

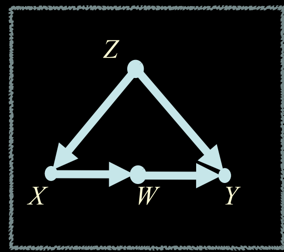

MOVING FROM THE “LAB” TO THE “REAL WORLD”...

Lab

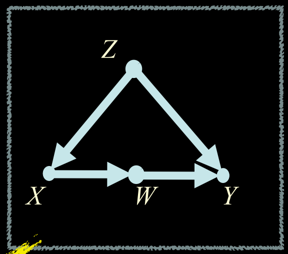

H_1

Real world

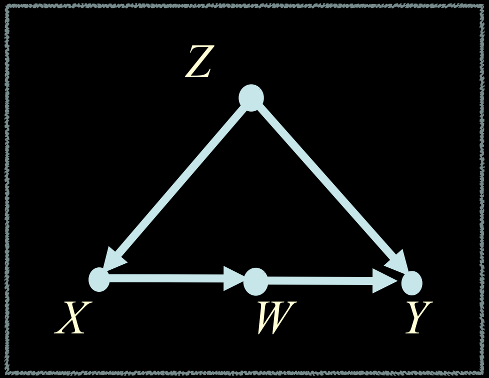

MOVING FROM THE “LAB” TO THE “REAL WORLD”...

Lab

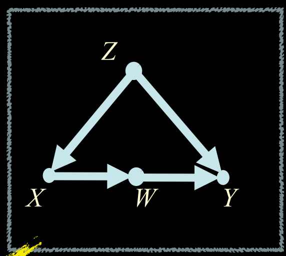

$f_z = f^*_z$
$f_x = f^*_x$
$f_w = f^*_w$
$f_y = f^*_y$

H_1

Real world

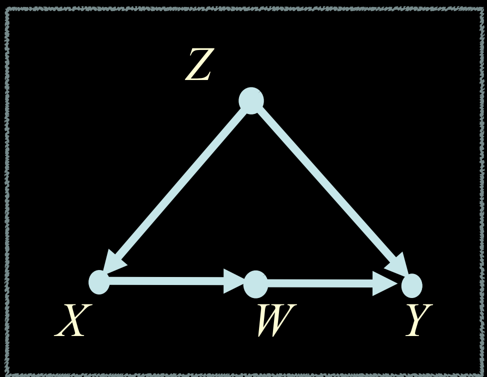

MOVING FROM THE "LAB" TO THE "REAL WORLD"...

Lab

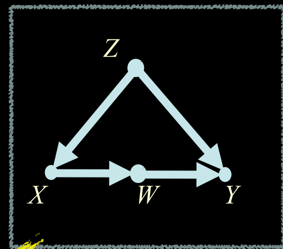

$$
\begin{aligned}
f_z &= f^*_z \\
f_x &= f^*_x \\
f_w &= f^*_w \\
f_y &= f^*_y
\end{aligned}
$$

H_1

Real world

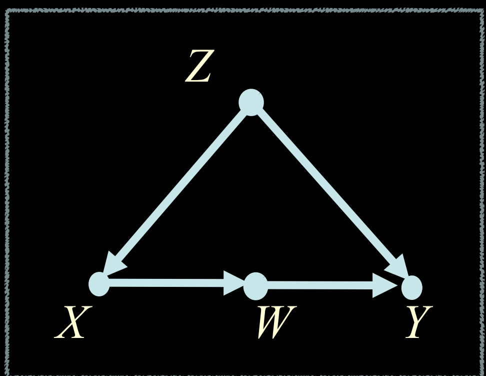

Everything is assumed to be
the same, trivially transportable!

MOVING FROM THE "LAB" TO THE "REAL WORLD"...

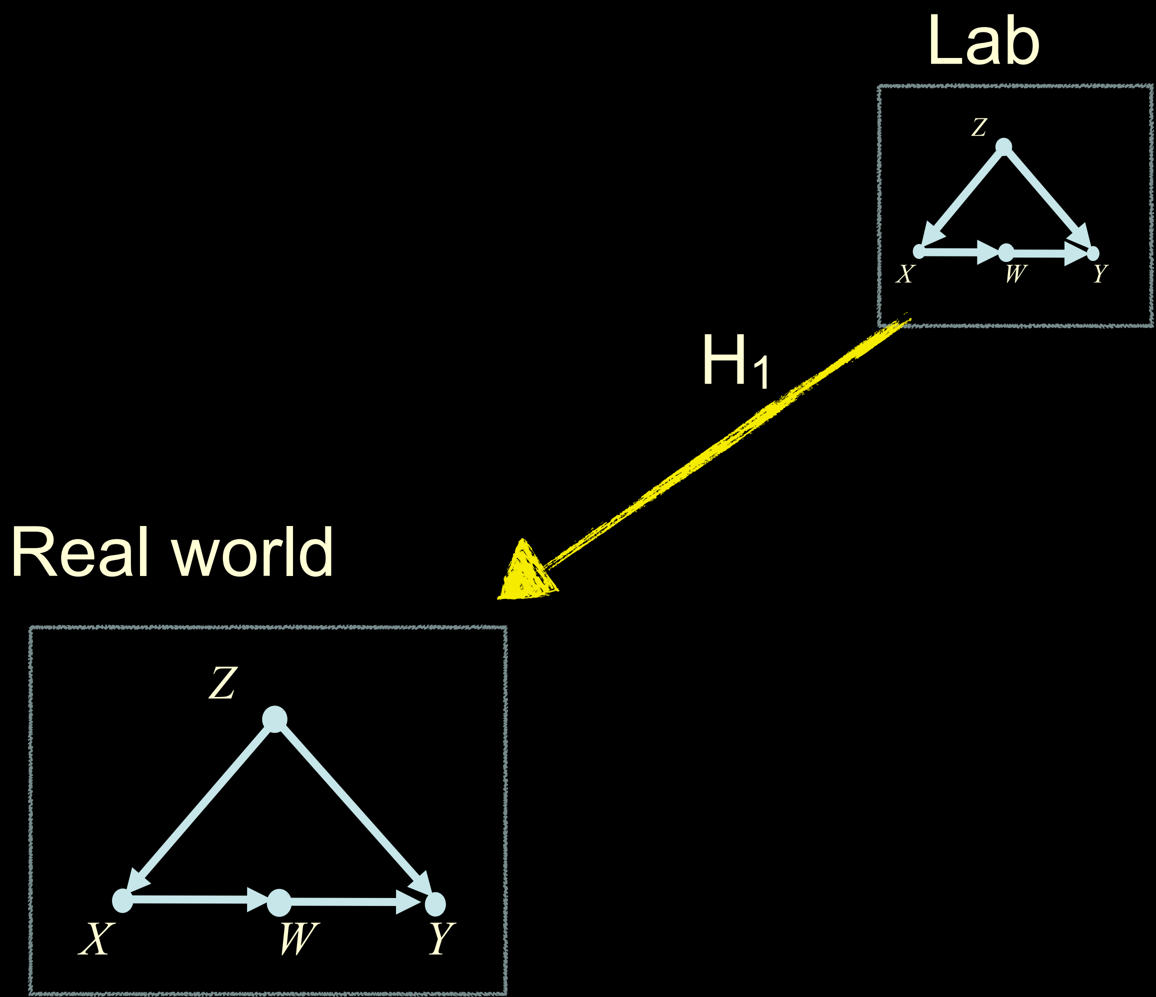

Everything is assumed to be
the same, trivially transportable!

MOVING FROM THE "LAB" TO THE "REAL WORLD"...

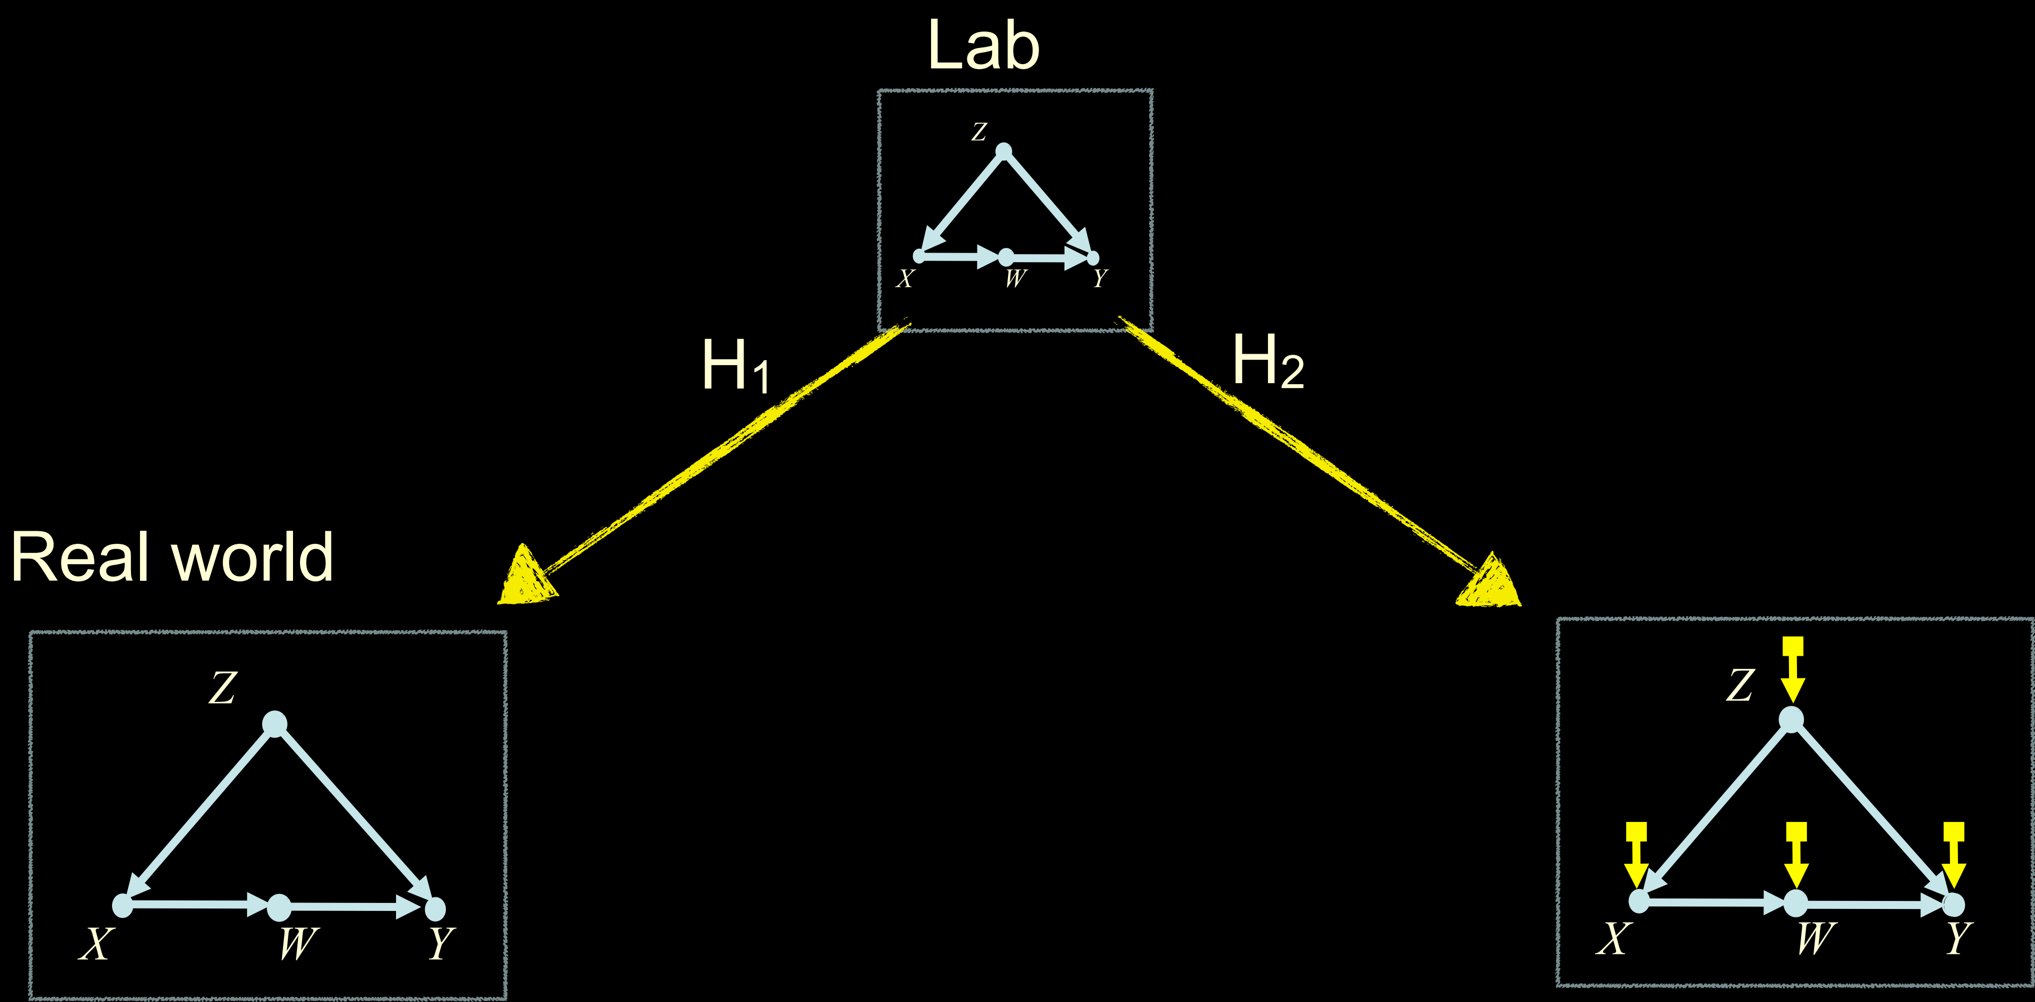

Everything is assumed to be
the same, trivially transportable!

MOVING FROM THE “LAB” TO THE “REAL WORLD”...

Lab

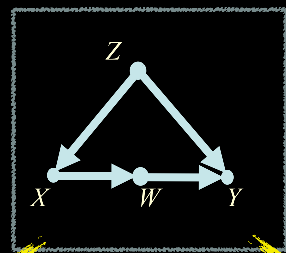

H_1 H_2

$f_Z \neq f^*_Z$
$f_X \neq f^*_X$
$f_W \neq f^*_W$
$f_Y \neq f^*_Y$

Real world

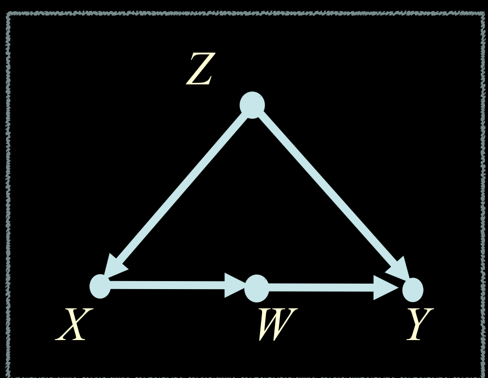

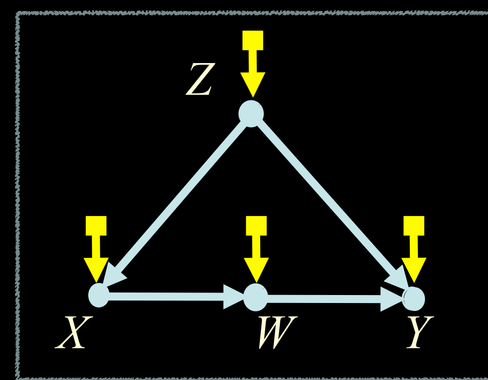

Everything is assumed to be
the same, trivially transportable!

MOVING FROM THE “LAB” TO THE “REAL WORLD”...

Lab

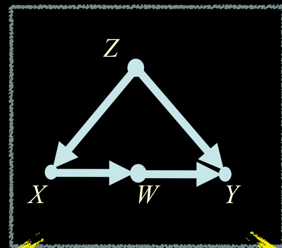

H_1

H_2

$f_Z \neq f^*_Z$
$f_X \neq f^*_X$
$f_W \neq f^*_W$
$f_Y \neq f^*_Y$

Real world

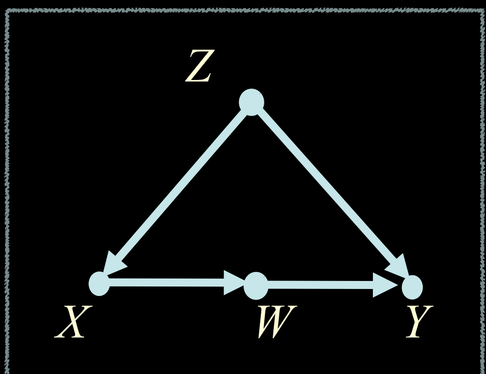

Everything is assumed to be the same, trivially transportable!

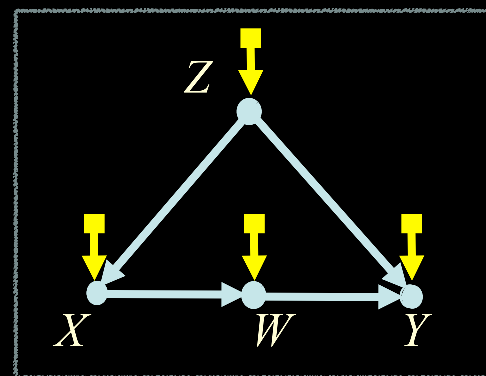

Everything is assumed to be different, not transportable...

MOVING FROM THE “LAB” TO THE “REAL WORLD”...

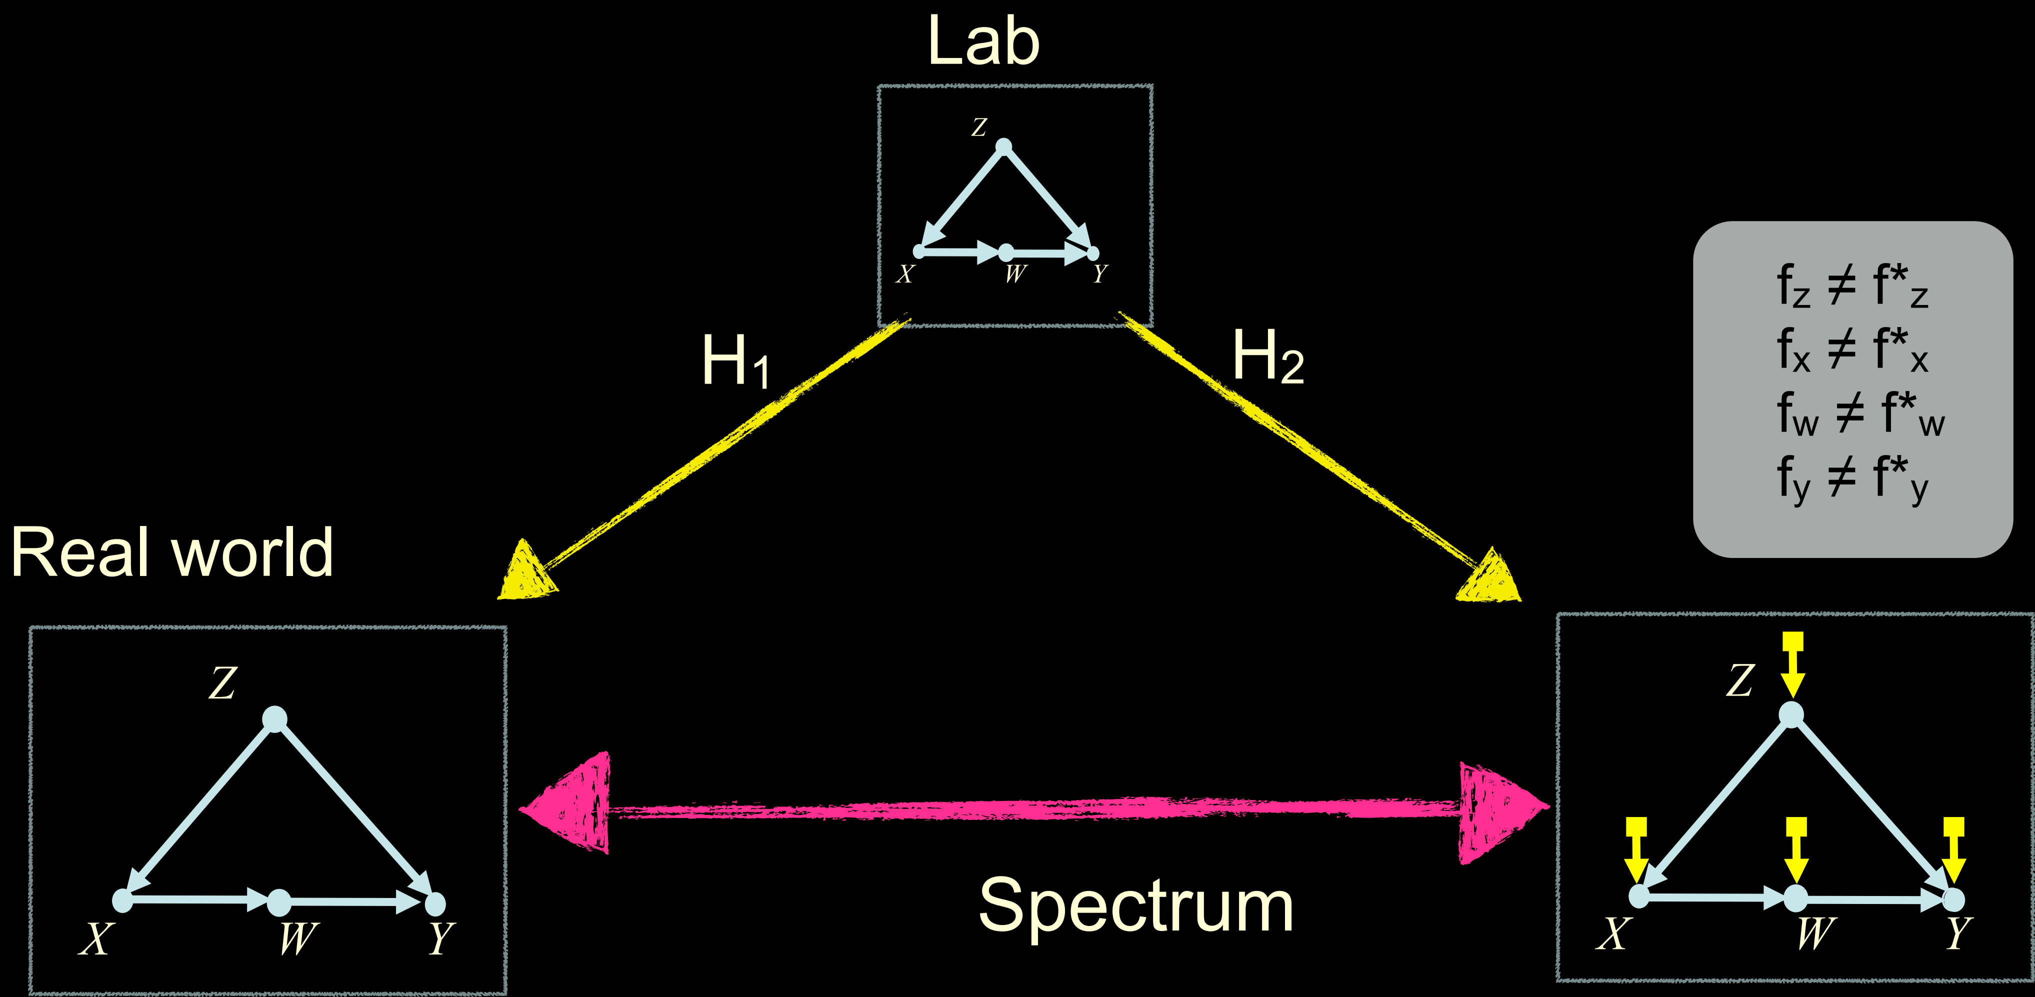

Everything is assumed to be the same, trivially transportable!

Everything is assumed to be different, not transportable...

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

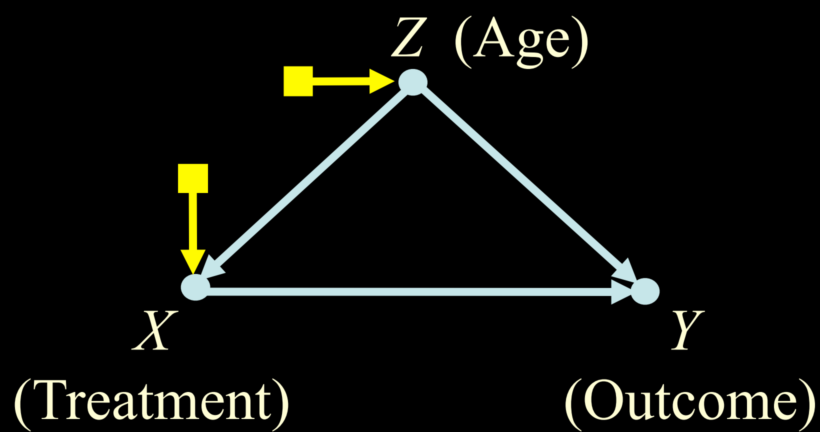

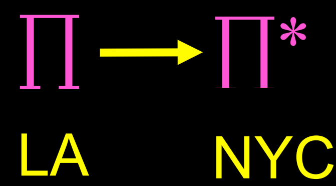

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

$f_z \neq f^*_z$
$f_x \neq f^*_x$
$f_y = f^*_y$

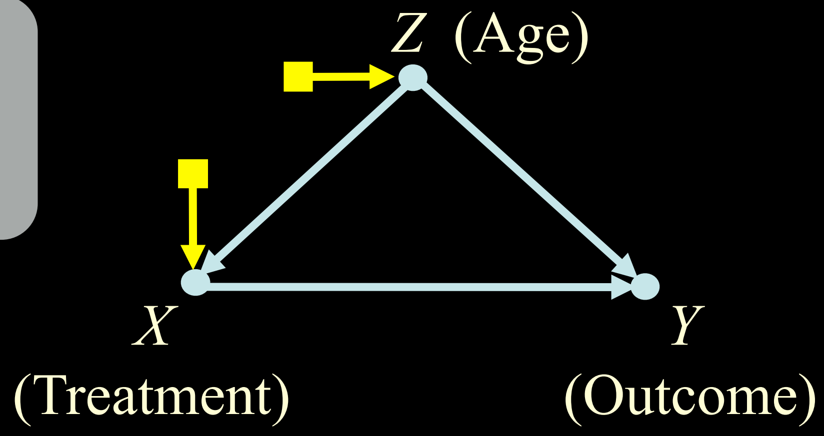

$\prod \longrightarrow \prod^*$

LA NYC

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

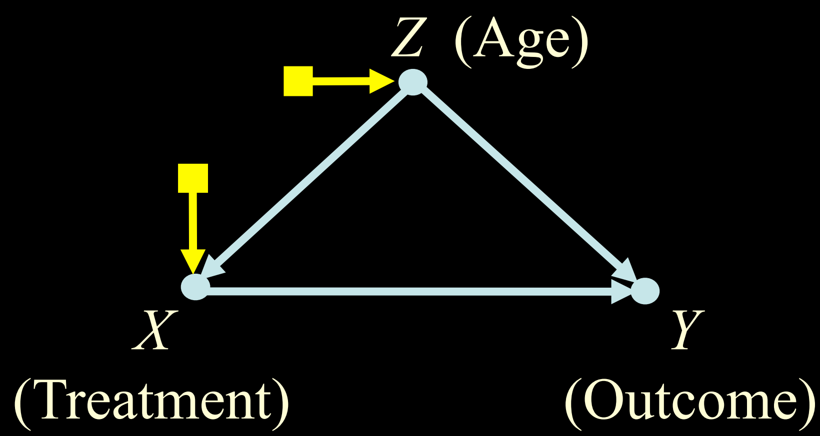

$\prod \longrightarrow \prod^*$

LA NYC

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

Input

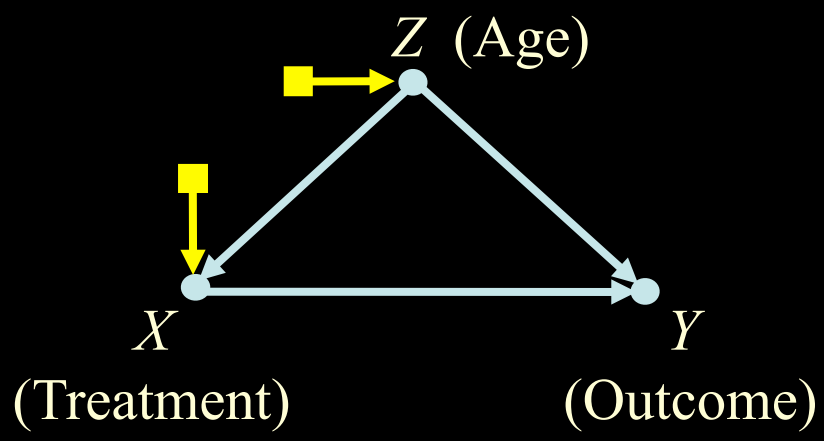

$\prod \longrightarrow \prod^*$

LA NYC

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

Input

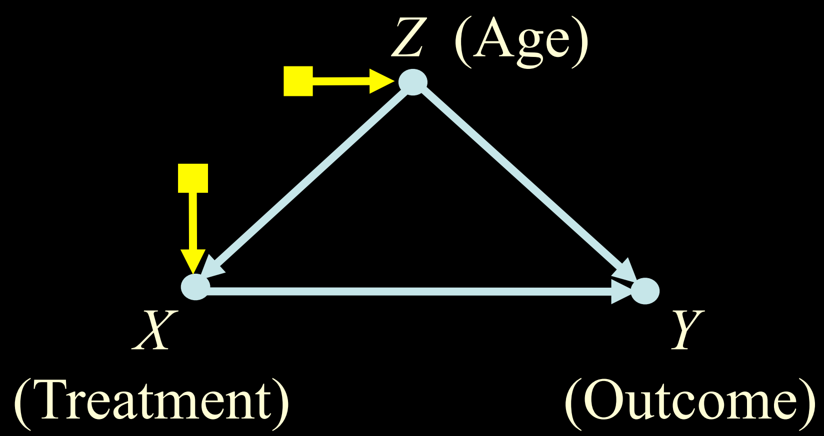

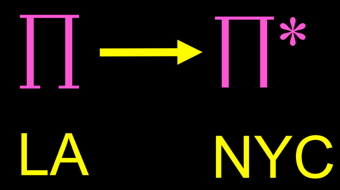

Experimental study in LA

Measured: $P(x, y, z)$

$P(y \mid do(x), z)$

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

Input

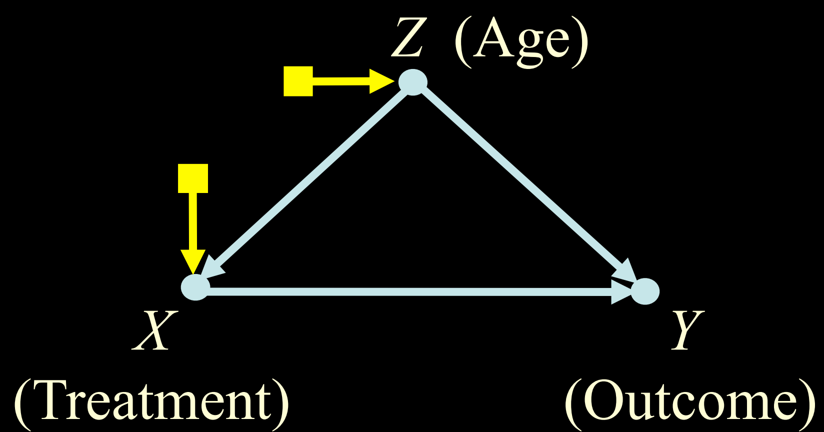

$\prod \longrightarrow \prod^*$

LA NYC

Experimental study in LA

Measured: $P(x, y, z)$

$P(y \mid do(x), z)$

Observational study in NYC

Measured: $P^*(x, y, z)$

$[\ P^*(z) \neq P(z)\ ]$

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

Input

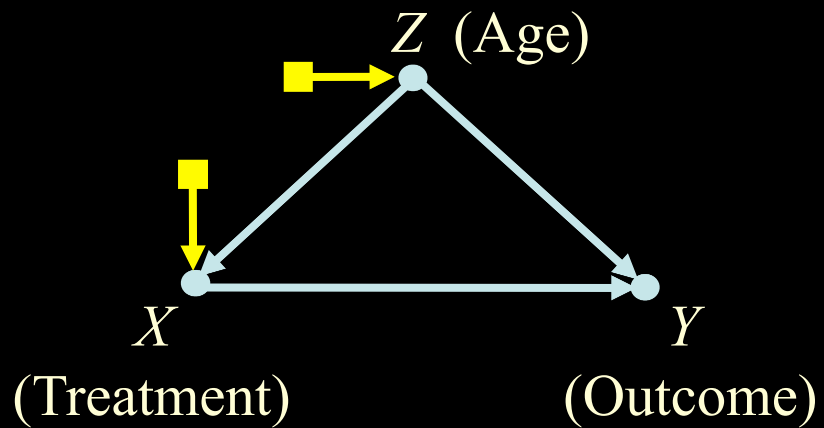

$\Pi \longrightarrow \Pi^*$

LA NYC

Experimental study in LA

Measured: $P(x, y, z)$

$P(y \mid do(x), z)$

Observational study in NYC

Measured: $P^*(x, y, z)$

$[\; P^*(z) \neq P(z) \;]$

Needed: Q = $P^*(y \mid do(x)) = ?$

MOTIVATION
WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

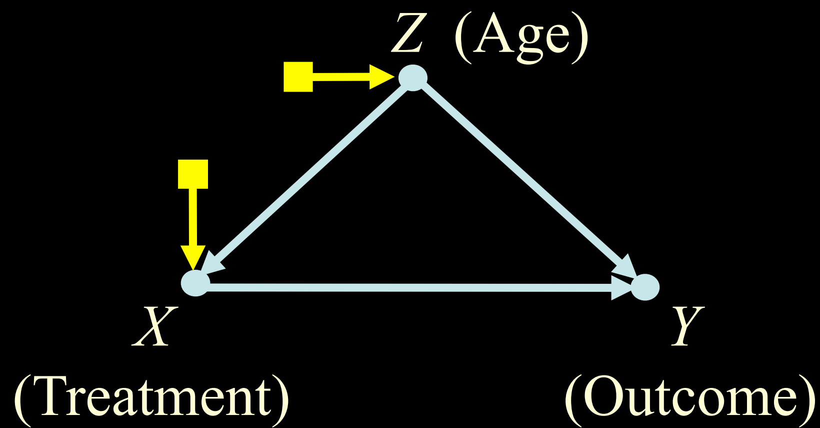

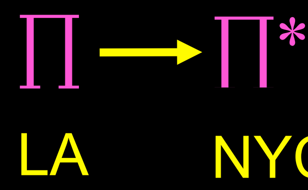

Input

Experimental study in LA

Measured: $P(x, y, z)$

$P(y \mid do(x), z)$

Observational study in NYC

Measured: $P^*(x, y, z)$

$[\ P^*(z) \neq P(z)\ ]$

Output

Needed: $Q = P^*(y \mid do(x)) = ?$

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

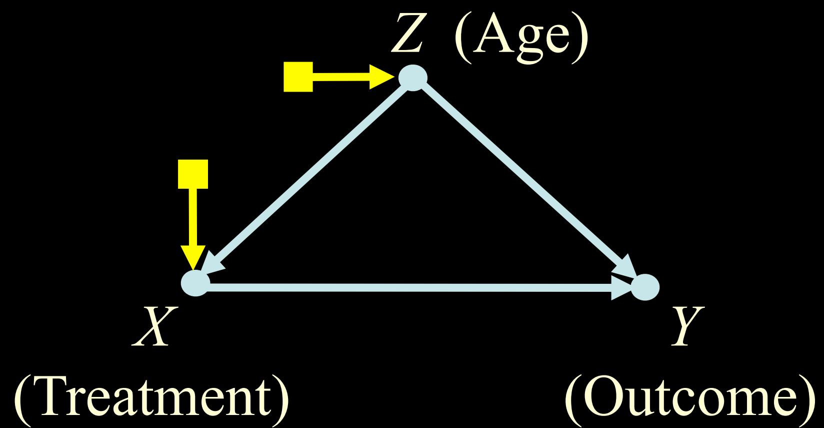

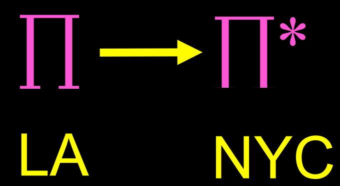

Experimental study in LA

Measured: $P(x, y, z)$

$P(y \mid do(x), z)$

Observational study in NYC

Measured: $P^*(x, y, z)$

$[\ P^*(z) \neq P(z)\ ]$

Needed: $Q = P^*(y \mid do(x)) = ?$

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

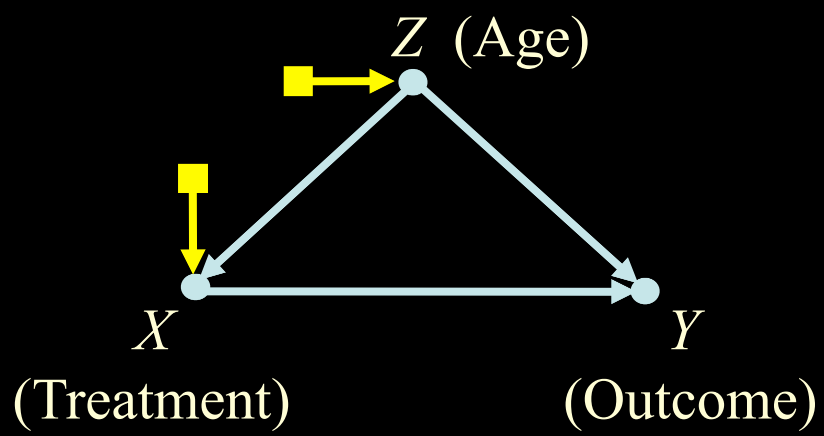

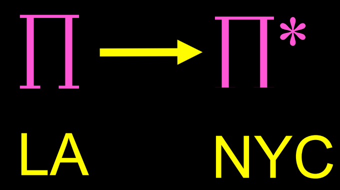

Experimental study in LA

Measured: $P(x, y, z)$

$P(y \mid do(x), z)$

Observational study in NYC

Measured: $P^*(x, y, z)$

$[\, P^*(z) \neq P(z) \,]$

Needed: $Q = P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z)\, P^*(z)$

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

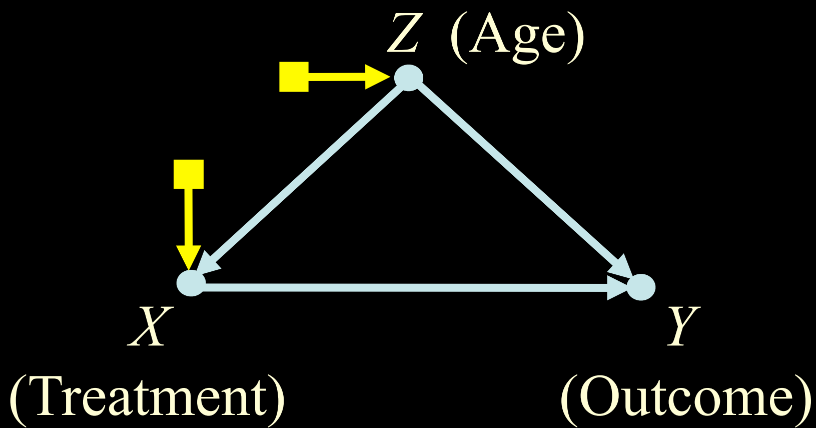

$$\underset{\text{LA}}{\prod} \longrightarrow \underset{\text{NYC}}{\prod^*}$$

Experimental study in LA

Measured: $P(x, y, z)$

$P(y \mid do(x), z)$

Observational study in NYC

Measured: $P^*(x, y, z)$

$[\ P^*(z) \neq P(z)\ ]$

Needed: $Q = P^*(y \mid do(x)) = \sum_Z \overset{\text{LA}}{P(y \mid do(x), z)}\ \overset{\text{NYC}}{P^*(z)}$

MOTIVATION

WHAT CAN EXPERIMENTS IN LA TELL US ABOUT NYC?

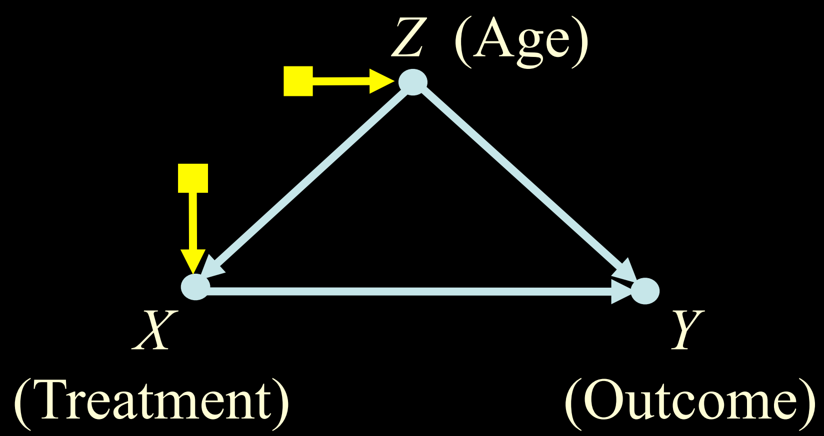

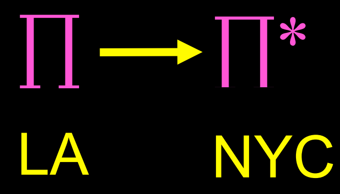

Experimental study in LA

Measured: $P(x, y, z)$

$P(y \mid do(x), z)$

Observational study in NYC

Measured: $P^*(x, y, z)$

$[\, P^*(z) \neq P(z) \,]$

Needed: $Q = P^*(y \mid do(x)) = \sum_z \underbrace{P(y \mid do(x), z)}_{\text{LA}} \underbrace{P^*(z)}_{\text{NYC}}$

Transport Formula (recalibration): $Q = F(P, P_{do}, P^*)$

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

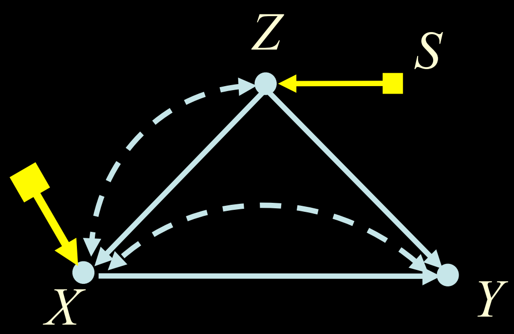

(a)

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

$z \leftarrow f_z(u_z, u_{xz})$
$x \leftarrow f_x(x, u_x, u_{xz}, u_{xy})$
$y \leftarrow f_y(x, z, u_y, u_{xy})$

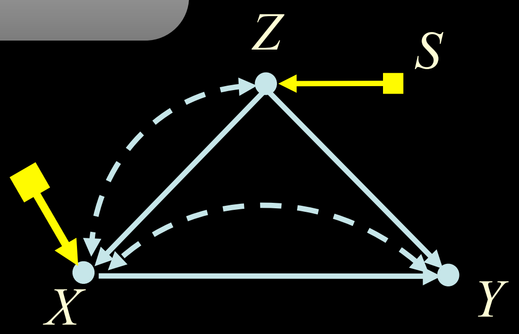

(a)

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

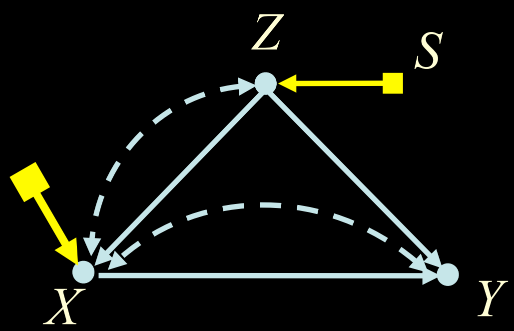

(a)

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

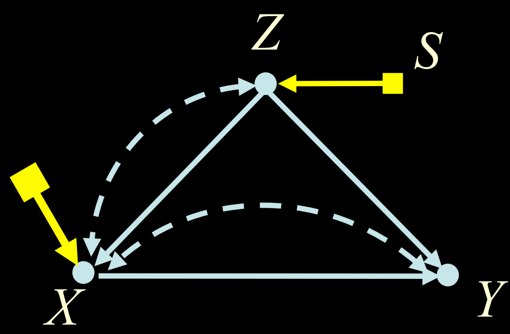

(a)

LA: P(y | do(x),z)

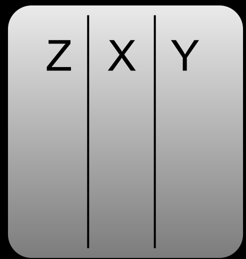

NYC: P*(x,y,z)

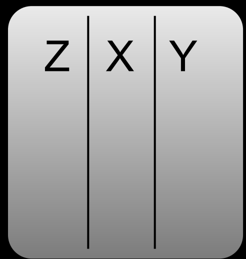

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

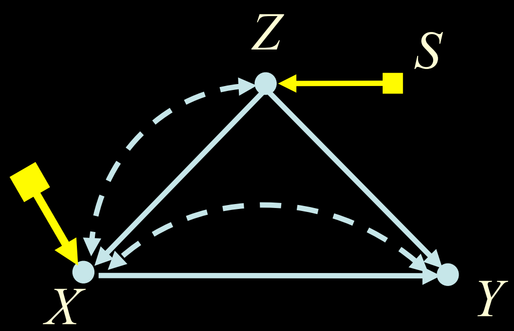

(a)

a) ***Z*** represents age

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z)$$

LA: P(y | do(x),z)

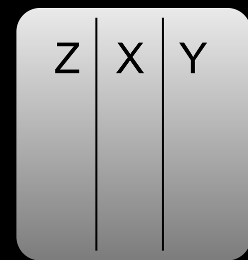

NYC: P*(x,y,z)

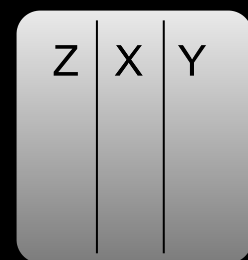

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

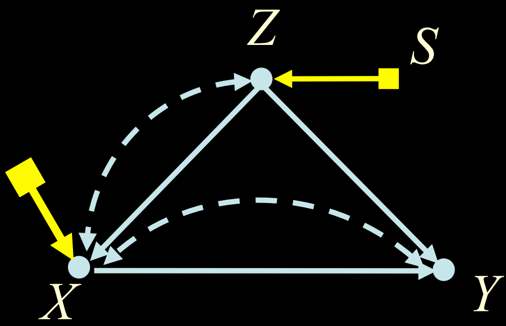

(a)

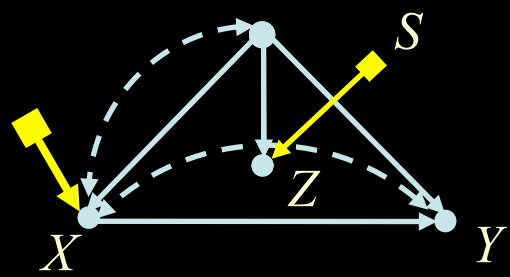

(b)

LA: P(y | do(x),z)

Z	X	Y

NYC: P*(x,y,z)

Z	X	Y

a) ***Z*** represents age

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z)$$

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

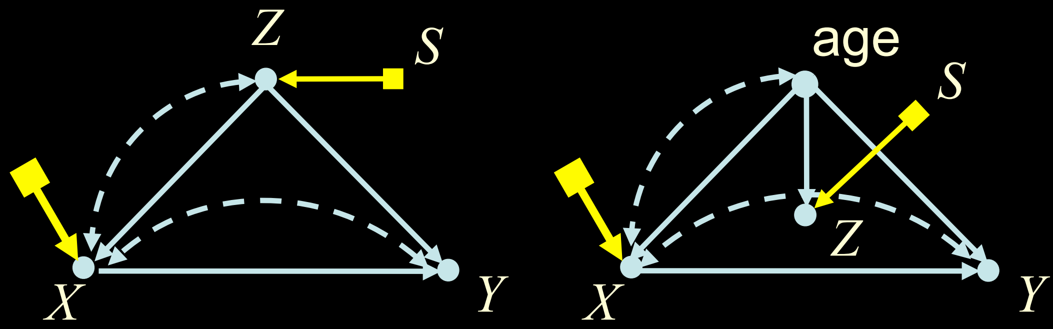

(a) (b)

LA: P(y | do(x),z)

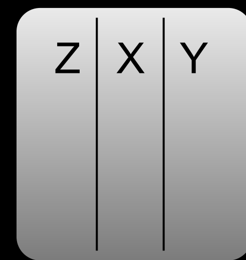

NYC: P*(x,y,z)

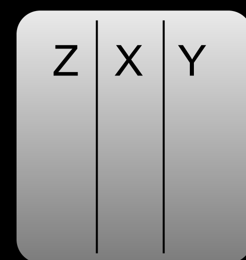

a) ***Z*** represents age

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z)$$

b) ***Z*** represents language skill

$$P^*(y \mid do(x)) = \ ?$$

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

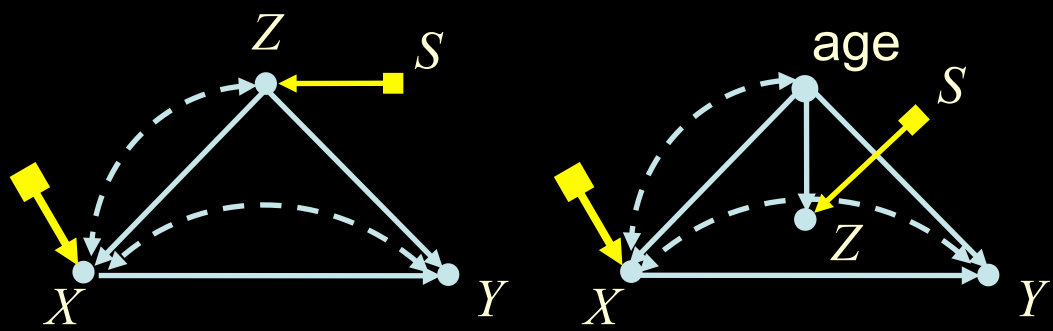

(a) (b)

a) $\boldsymbol{Z}$ represents age

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z)$$

b) $\boldsymbol{Z}$ represents language skill

$$P^*(y \mid do(x)) = P(y \mid do(x))$$

LA: P(y | do(x),z)

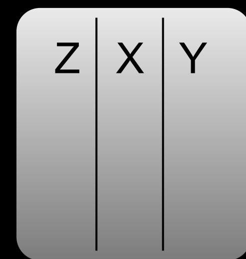

NYC: P*(x,y,z)

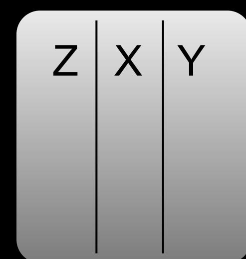

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

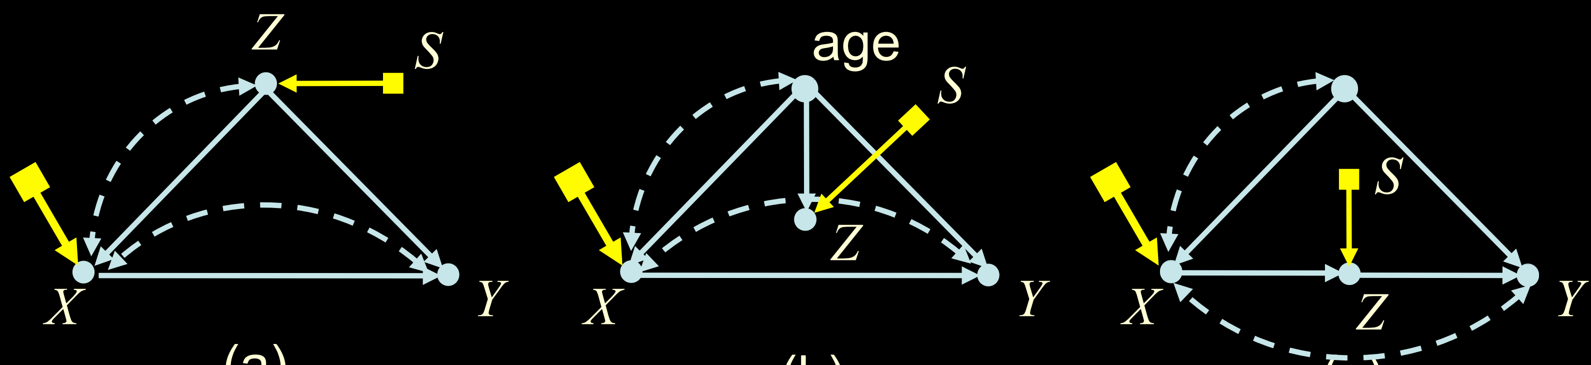

(a) (b) (c)

LA: P(y | do(x),z)

Z	X	Y

NYC: P*(x,y,z)

Z	X	Y

a) Z represents age

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z)$$

b) Z represents language skill

$$P^*(y \mid do(x)) = P(y \mid do(x))$$

c) Z represents a bio-marker

$$P^*(y \mid do(x)) = ?$$

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

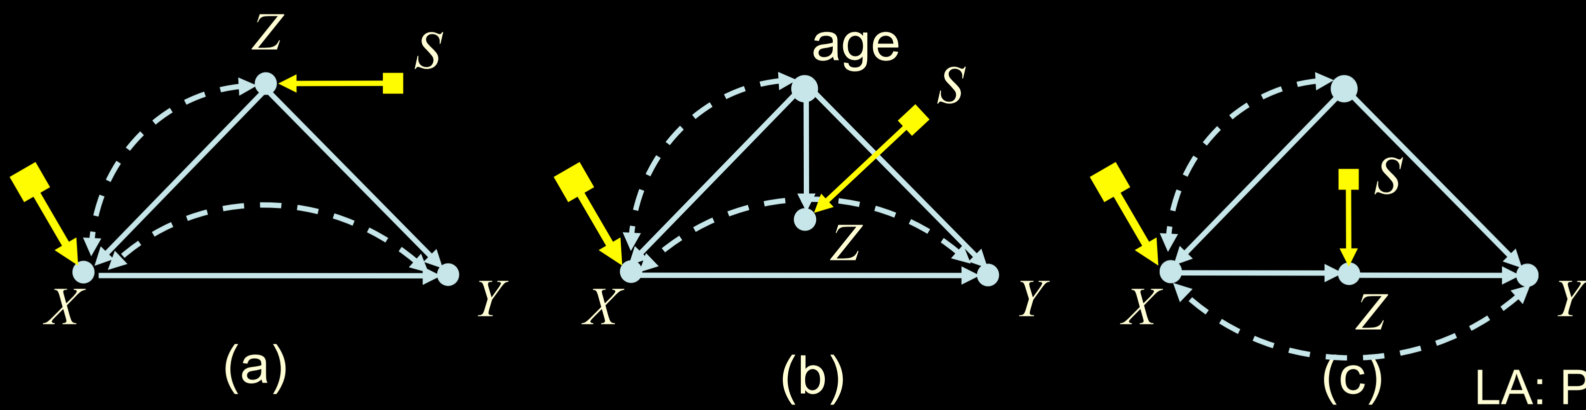

LA: P(y | do(x),z)

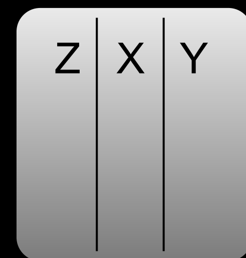

a) ***Z*** represents age

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z)$$

b) ***Z*** represents language skill

$$P^*(y \mid do(x)) = P(y \mid do(x))$$

NYC: P*(x,y,z)

c) ***Z*** represents a bio-marker

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z \mid x)$$

TRANSPORT FORMULA SENSITIVITY TO THE CAUSAL ASSUMPTIONS

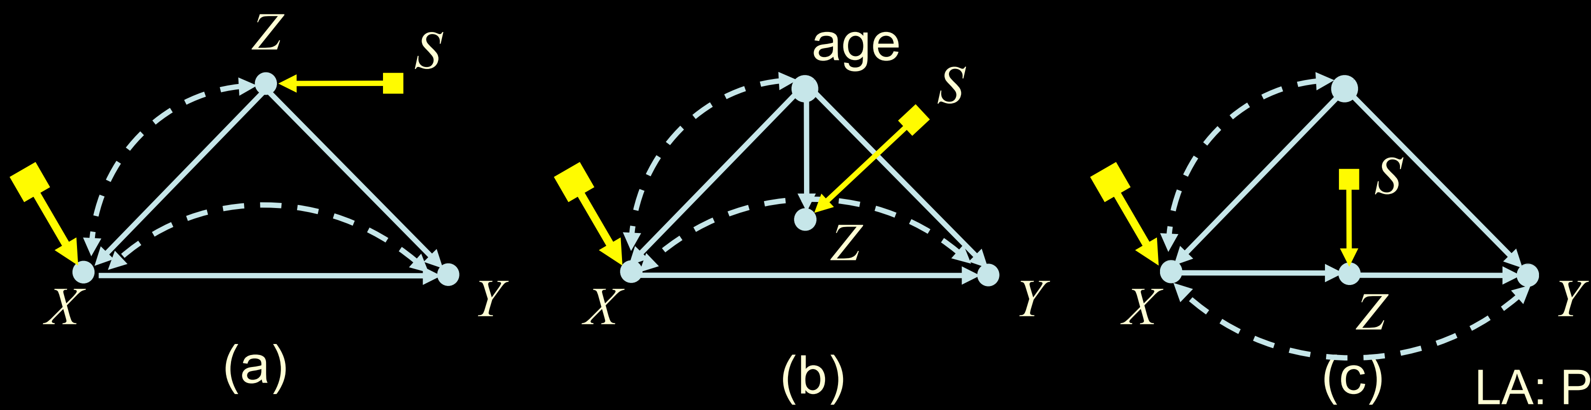

LA: P(y | do(x),z)

Z	X	Y

a) $\boldsymbol{Z}$ represents age

Lesson. Causal assumptions are required since the data does not impose enough constraints over the causal structure, and the results are structure-sensitive.

NYC: P*(x,y,z)

Z	X	Y

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z \mid x)$$

DECISION PROBLEM. IS THE TARGET EFFECT TRANSPORTABLE?

1 **Query**

Q= $P^*(y \mid do(x))$

DECISION PROBLEM.
IS THE TARGET EFFECT TRANSPORTABLE?

1 **Query**

Q= $P^*(y \mid do(x))$

2 **Graph**

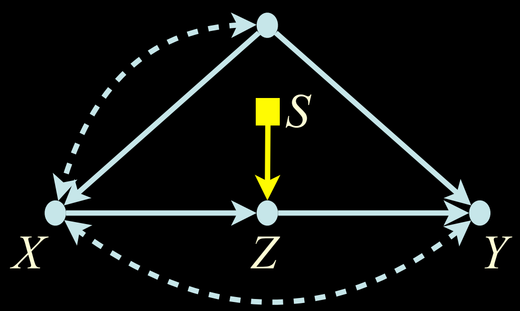

DECISION PROBLEM.
IS THE TARGET EFFECT TRANSPORTABLE?

1 **Query**

Q= $P^*(y \mid do(x))$

2 **Graph**

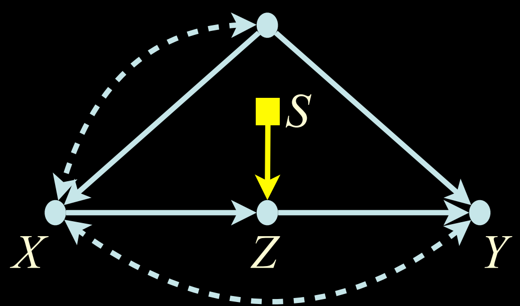

3 **Data**

$$P(x, z, y), P(z, y \mid do(x)), P^*(x, z, y)$$

DECISION PROBLEM.
IS THE TARGET EFFECT TRANSPORTABLE?

Q: Based on the current knowledge about the problem (2) and the available datasets (3), is the research question (1) transportable?

1 Query

Q= $P^*(y \mid do(x))$

2 Graph

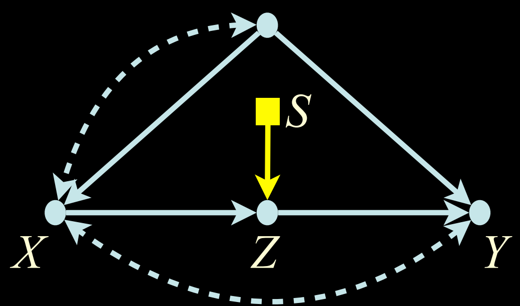

3 Data

$$P(x, z, y), P(z, y \mid do(x)), P^*(x, z, y)$$

DECISION PROBLEM. IS THE TARGET EFFECT TRANSPORTABLE?

Q: Based on the current knowledge about the problem (2) and the available datasets (3), is the research question (1) transportable?

1 **Query**

Q= $P^*(y \mid do(x))$

2 **Graph**

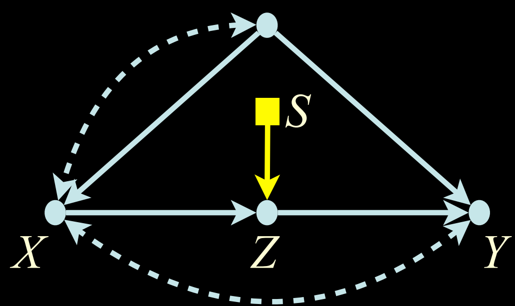

3 **Data**

$P(x, z, y), P(z, y \mid do(x)), P^*(x, z, y)$

TR Inference Engine

DECISION PROBLEM.
IS THE TARGET EFFECT TRANSPORTABLE?

Q: Based on the current knowledge about the problem (2) and the available datasets (3), is the research question (1) transportable?

1 Query

Q= $P^*(y \mid do(x))$

2 Graph

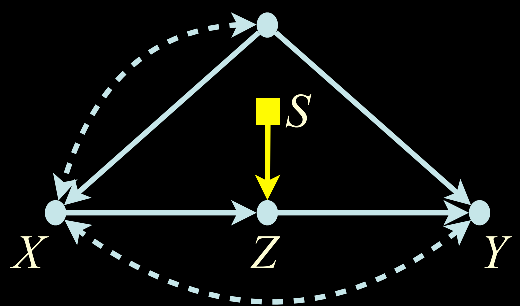

3 Data

$P(x, z, y), P(z, y \mid do(x)), P^*(x, z, y)$

TR Inference Engine → **Solution**

DECISION PROBLEM.
IS THE TARGET EFFECT TRANSPORTABLE?

Q: Based on the current knowledge about the problem (2) and the available datasets (3), is the research question (1) transportable?

1 **Query**

Q= $P^*(y \mid do(x))$

2 **Graph**

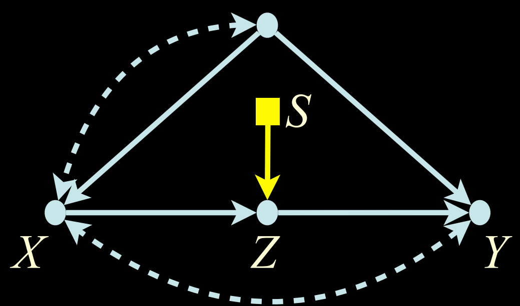

3 **Data**

$P(x, z, y), P(z, y \mid do(x)), P^*(x, z, y)$

TR Inference Engine

Solution

yes / no

DECISION PROBLEM.
IS THE TARGET EFFECT TRANSPORTABLE?

Q: Based on the current knowledge about the problem (2) and the available datasets (3), is the research question (1) transportable?

1 **Query**

Q= $P^*(y \mid do(x))$

2 **Graph**

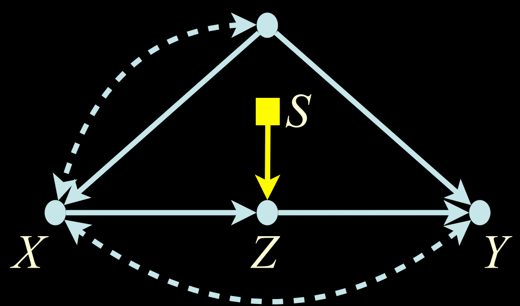

3 **Data**

$$P(x, z, y), P(z, y \mid do(x)), P^*(x, z, y)$$

TR Inference Engine

Solution

yes / no

DECISION PROBLEM. IS THE TARGET EFFECT TRANSPORTABLE?

Q: Based on the current knowledge about the problem (2) and the available datasets (3), is the research question (1) transportable?

1 **Query**

Q= $P^*(y \mid do(x))$

2 **Graph**

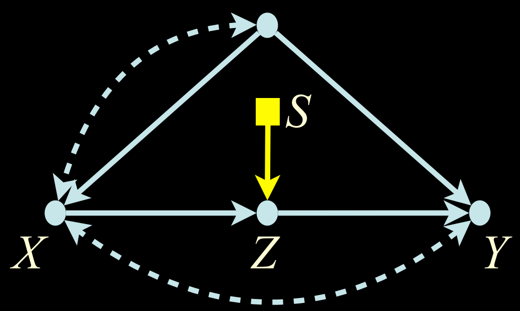

3 **Data**

$P(x, z, y), P(z, y \mid do(x)), P^*(x, z, y)$

TR Inference Engine

Solution

yes / no

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) P^*(z \mid x)$$

Causation in target domain (Query)

Obs. Data in target + Exp. Data in source (Data)

👍

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x))$

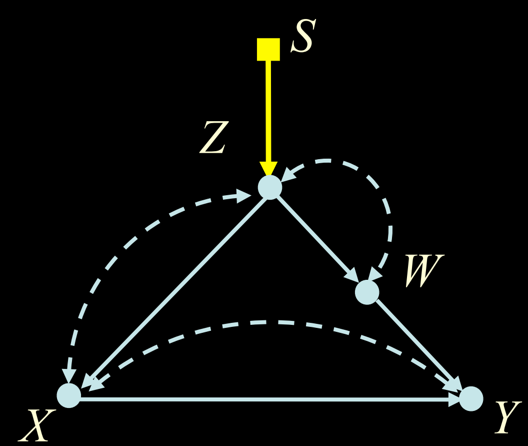

G

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x))$ query

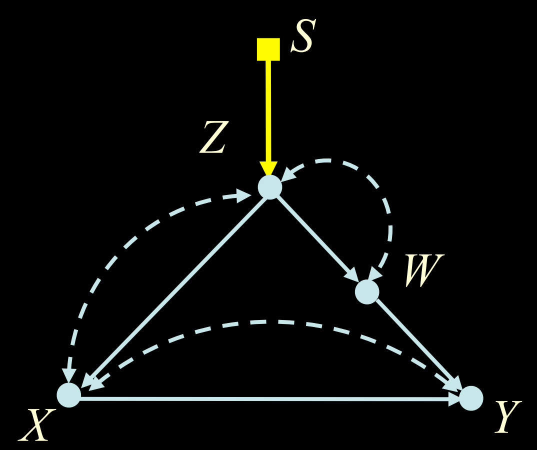

G

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity $\boldsymbol{Q}$ is transportable from Π to Π^* (G) **if and only if** there exists a do-calculus reduction of $\boldsymbol{Q}(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x))$ query

$= P(y \mid do(x), s)$

$= \sum_w P(y \mid do(x), s, w)\, P(w \mid do(x), s)$

$= \sum_w P(y \mid do(x), w)\, P(w \mid do(x), s)$

$= \sum_w P(y \mid do(x), w)\, P(w \mid s)$

$= \sum_w P(y \mid do(x), w)\, P^*(w)$

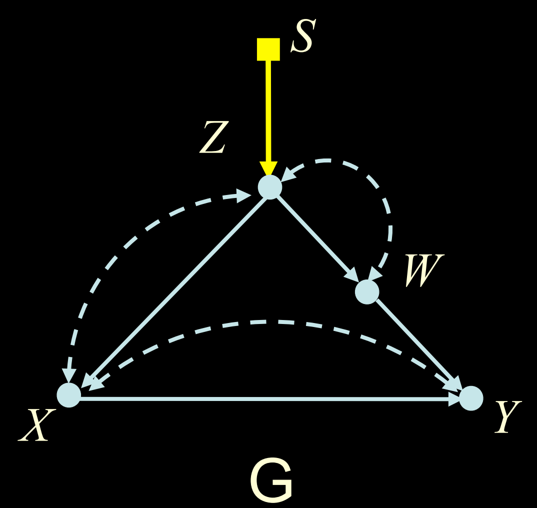

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x))$ query

$= P(y \mid do(x), s)$

$= \sum_w P(y \mid do(x), s, w)\, P(w \mid do(x), s)$

$= \sum_w P(y \mid do(x), w)\, P(w \mid do(x), s)$

$= \sum_w P(y \mid do(x), w)\, P(w \mid s)$

$= \sum_w P(y \mid do(x), w)\, P^*(w)$ 👍

data

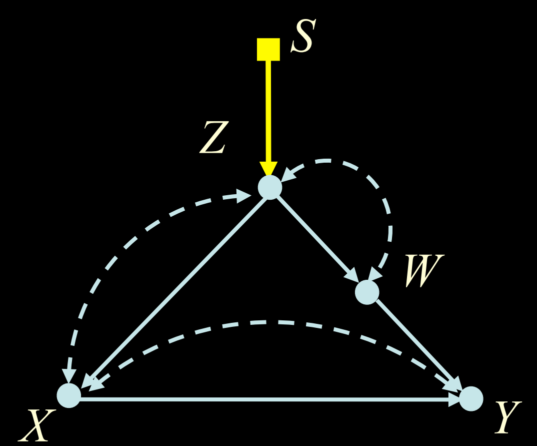

G

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

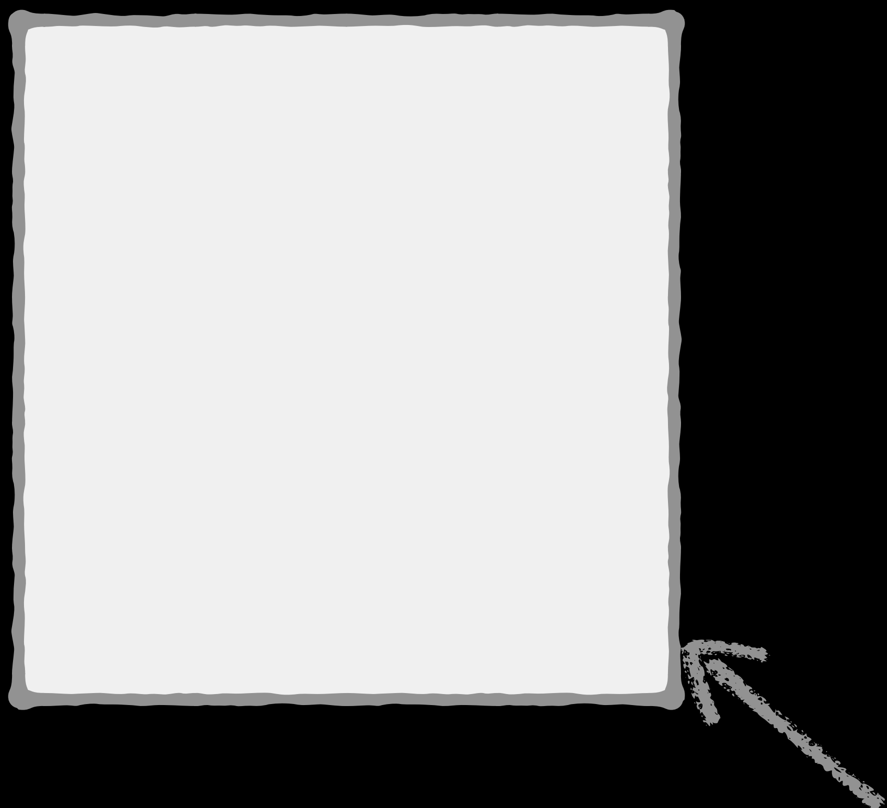

All models

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is
transportable in G

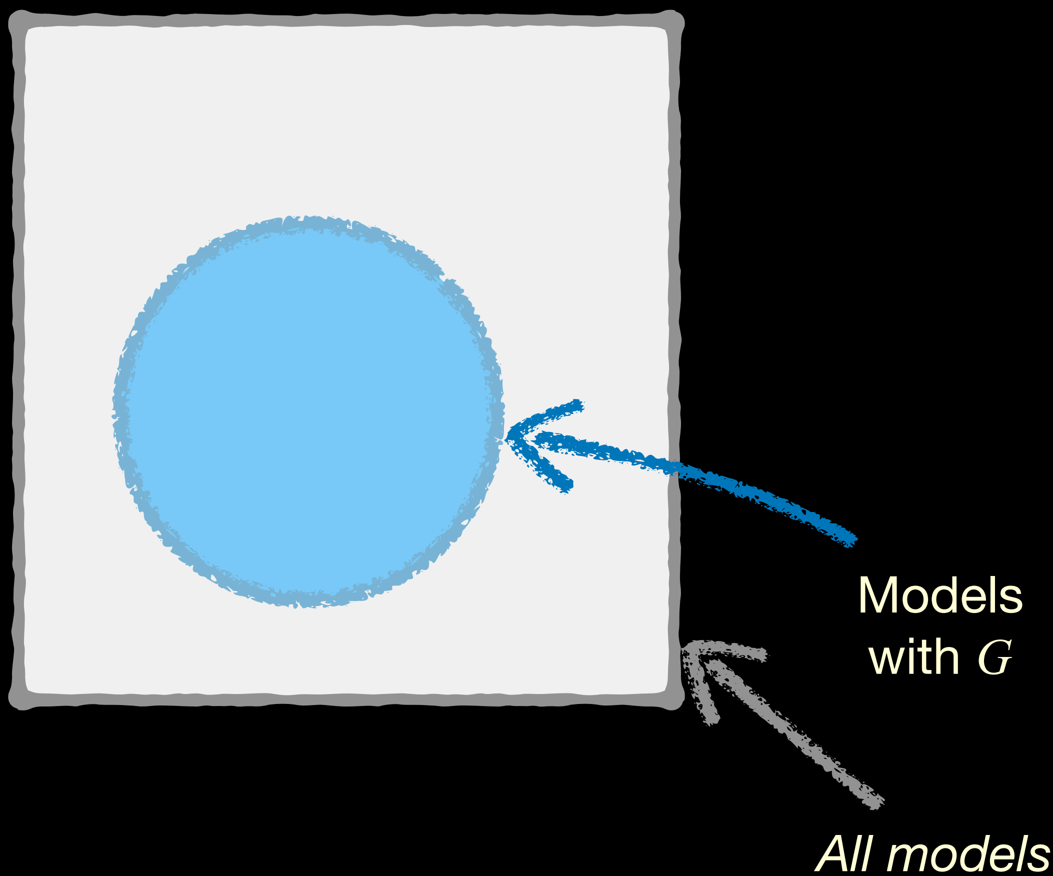

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

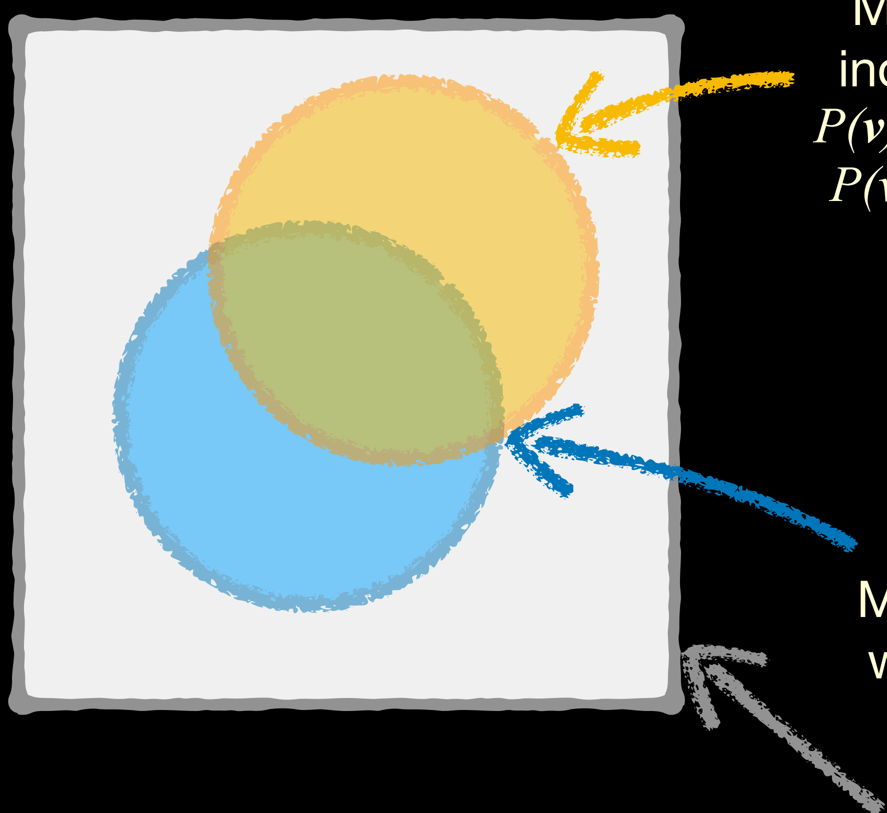

Models inducing $P(\mathbf{v}), P^*(\mathbf{v}), P(v|do(x))$

Models with G

All models

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

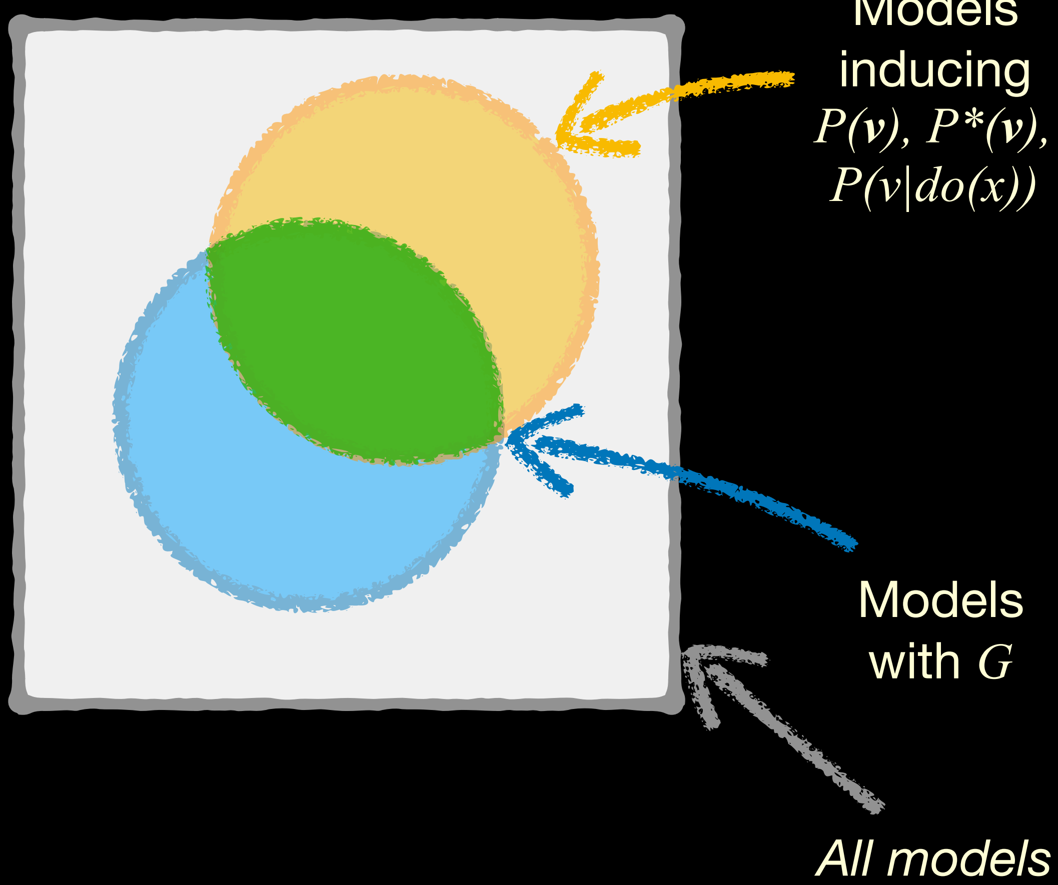

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

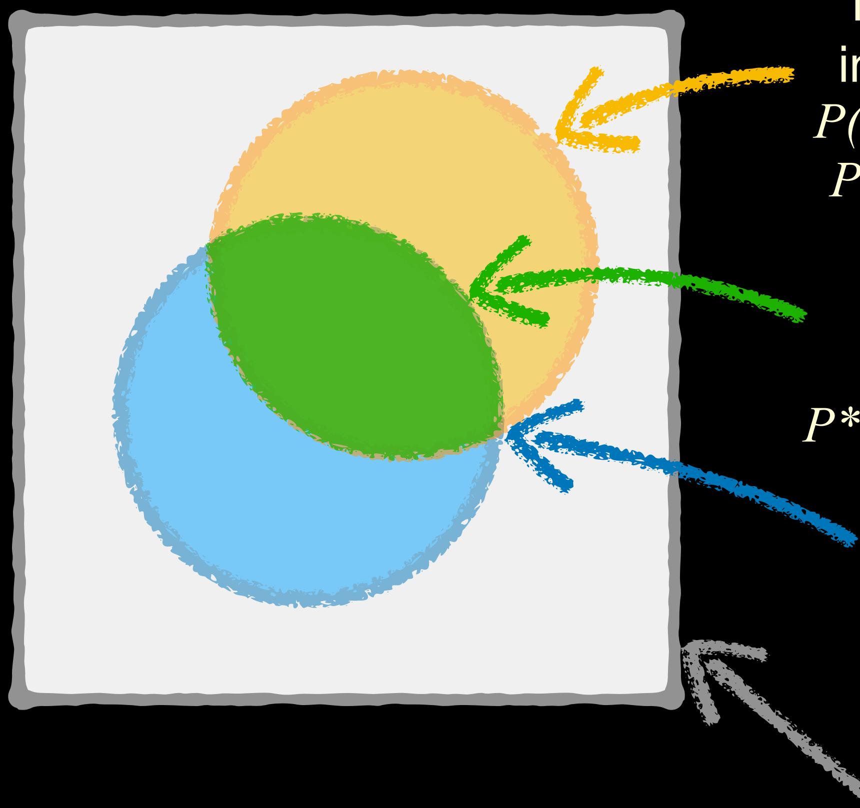

MEANING OF TRANSPORTABILITY

P(y|do(x))* is transportable in *G*

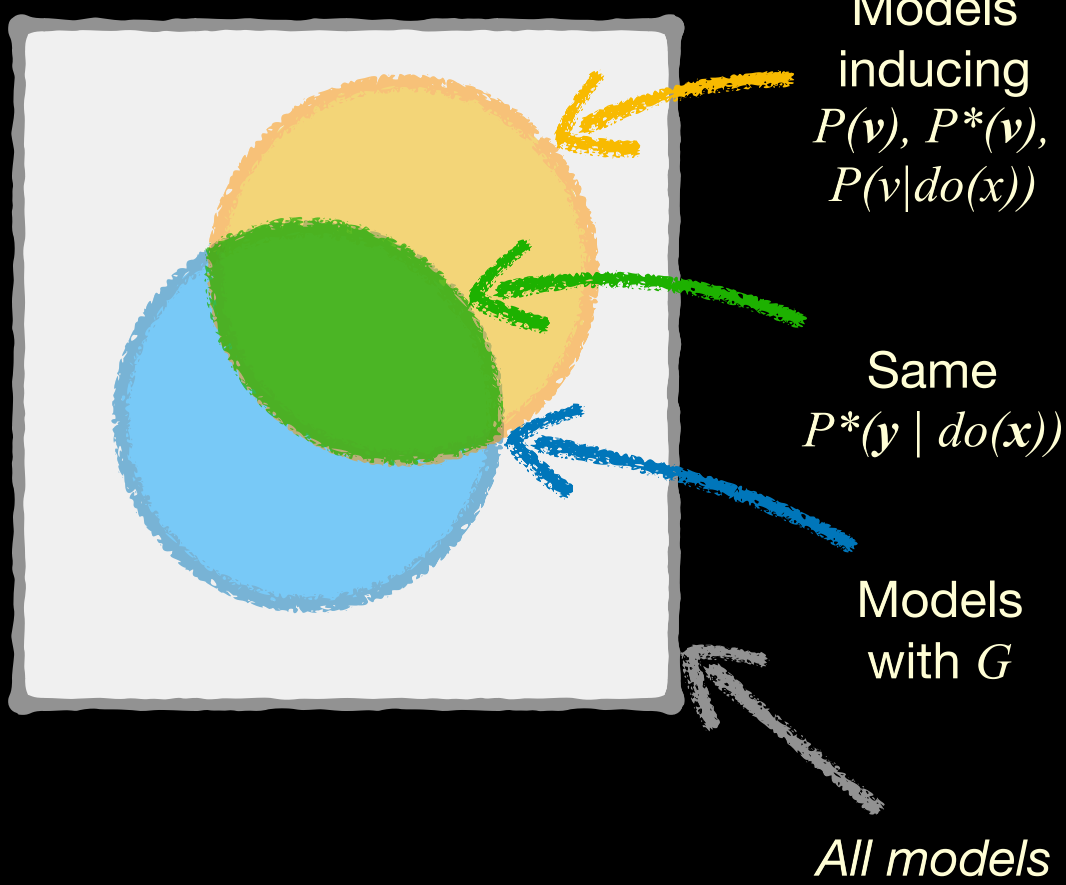

Formally,
for any two causal models (encoding the unobserved nature),
N_1, N_2, s.t.

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

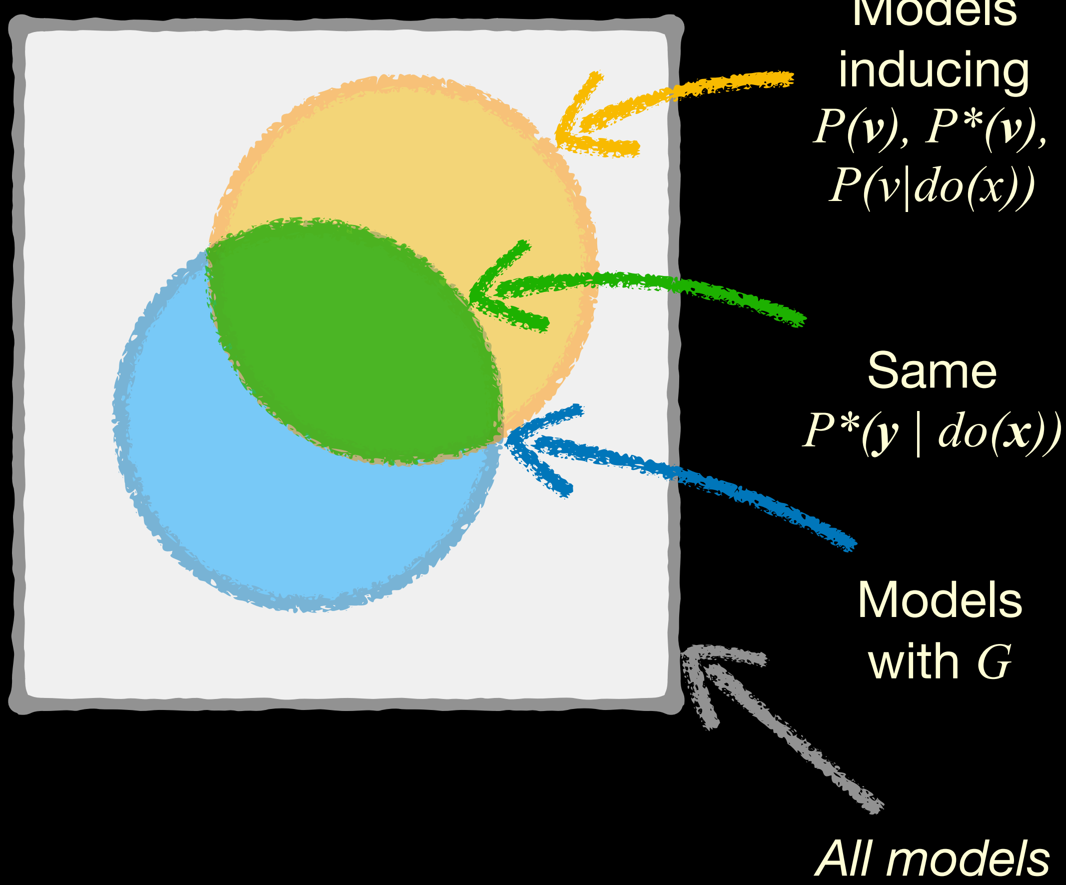

Formally, for any two causal models (encoding the unobserved nature), N_1, N_2, s.t. $G(N_1) = G(N_2) = G$,

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

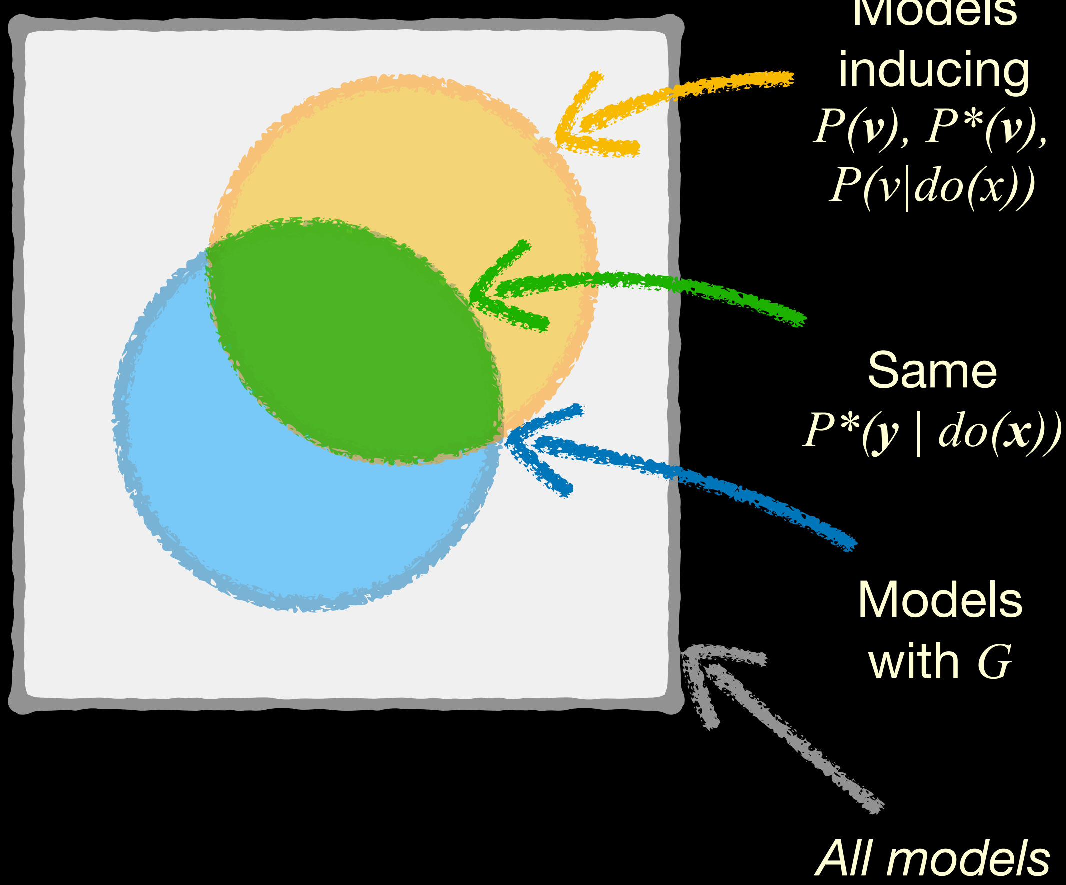

Formally, for any two causal models (encoding the unobserved nature), N_1, N_2, s.t.

$G(N_1) = G(N_2) = G$,

$P_1(v) = P_2(v)$,

$P^*_1(v) = P^*_2(v)$,

$P_1(v \mid do(x)) = P_2(v \mid do(x))$

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

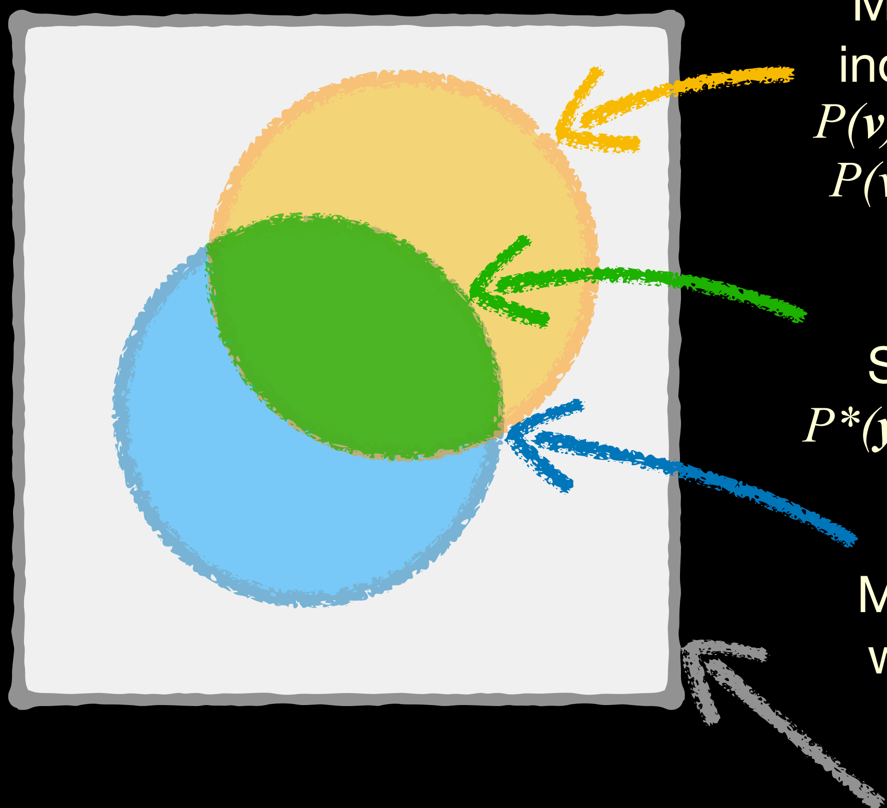

Formally,
for any two causal models (encoding the unobserved nature),
N_1, N_2, s.t.

$G(N_1) = G(N_2) = G$,

$P_1(v) = P_2(v)$,
$P^*_1(v) = P^*_2(v)$,
$P_1(v \mid do(x)) = P_2(v \mid do(x))$, then

$P^*_1(y \mid do(x)) = P^*_2(y \mid do(x))$.

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

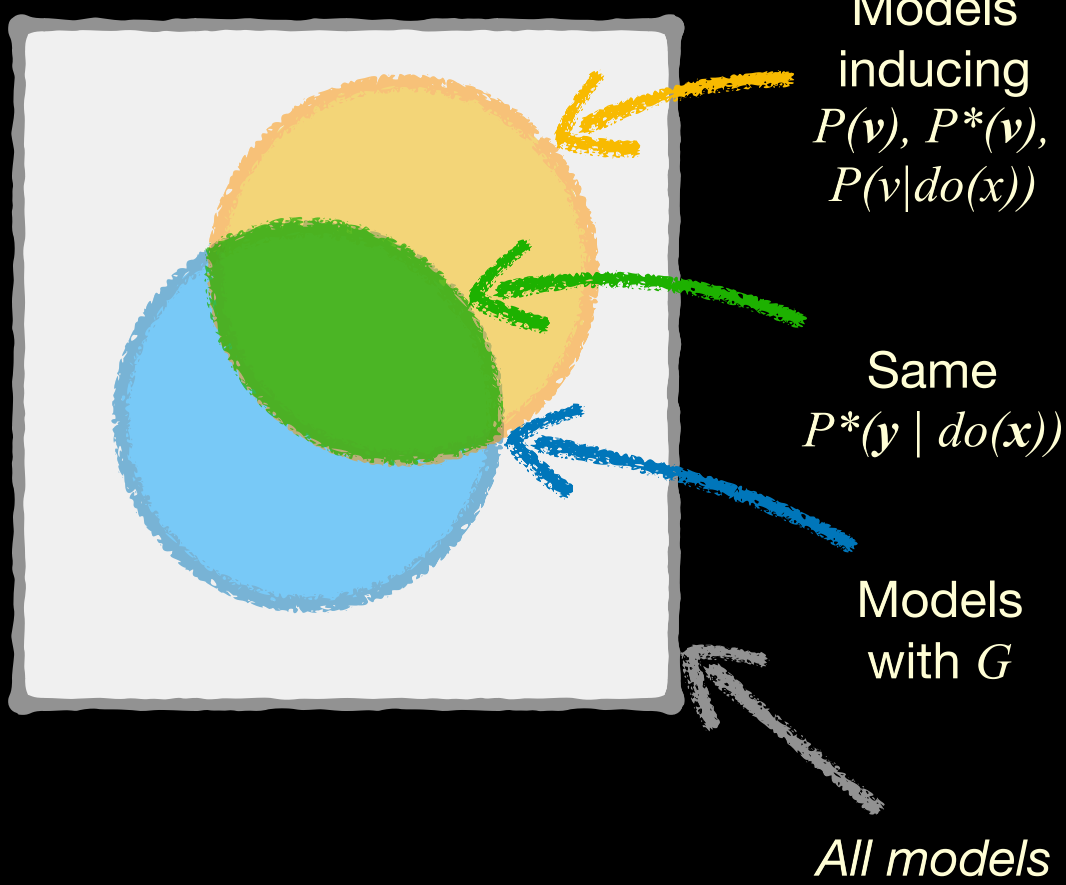

Formally, for any two causal models (encoding the unobserved nature), N_1, N_2, s.t.

$G(N_1) = G(N_2) = G$,

$P_1(v) = P_2(v)$,

$P^*_1(v) = P^*_2(v)$,

$P_1(v \mid do(x)) = P_2(v \mid do(x))$, then

$P^*_1(y \mid do(x)) = P^*_2(y \mid do(x))$.

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x)) = ?$

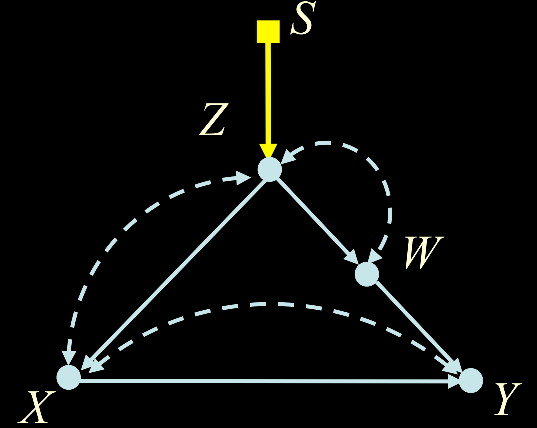

G

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) **if and only if** there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x)) = ?$

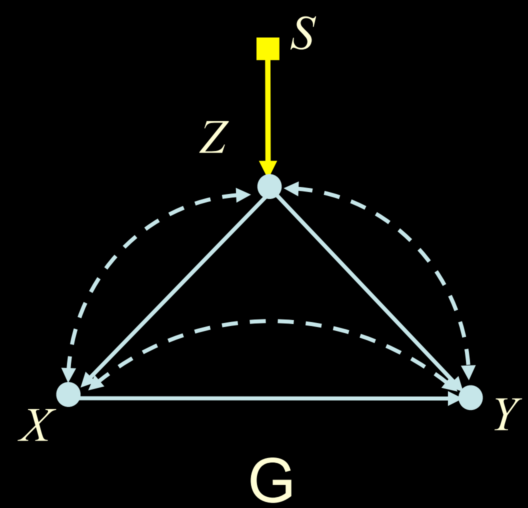

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity $\boldsymbol{Q}$ is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $\boldsymbol{Q(\Pi^*)}$ to an estimand that is a function of the observed distributions.

Q = P*(y | do(x)) = ?

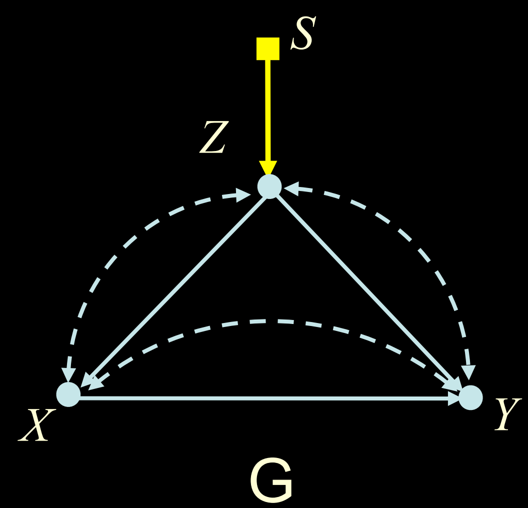

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity $\boldsymbol{Q}$ is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $\boldsymbol{Q}(\Pi^*)$ to an estimand that is a function of the observed distributions.

Q = P*(y | do(x)) = ? Nope.

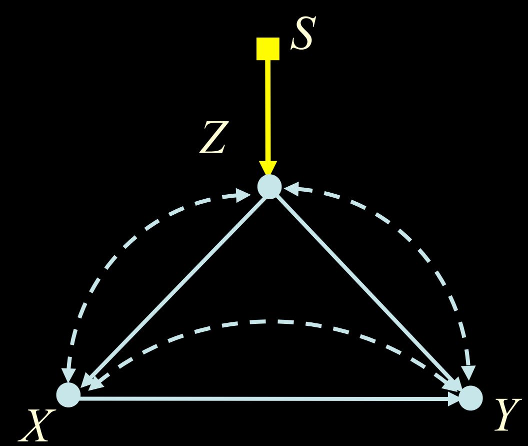

G

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity $\boldsymbol{Q}$ is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $\boldsymbol{Q}(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x)) = ?$ Nope.

Proof. $\exists$ N_1, N_2

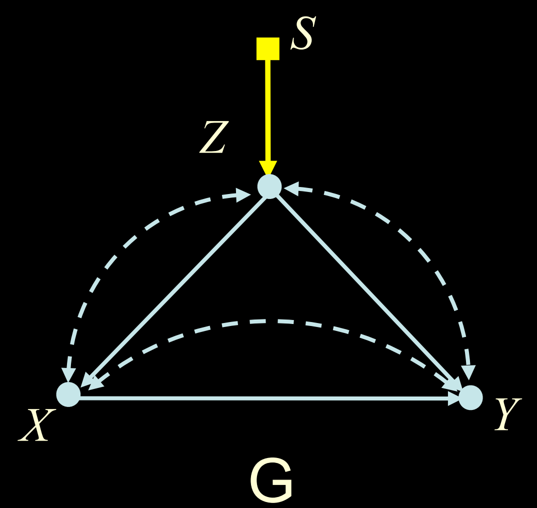

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity $\boldsymbol{Q}$ is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $\boldsymbol{Q}(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x)) = ?$ Nope.

Proof. ∃ N_1, N_2

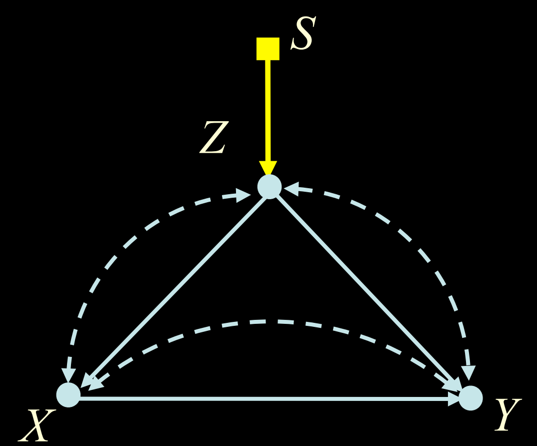

G

N_1

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow U_y \end{cases}$$

$P(u_i) = 1/2$

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity $\boldsymbol{Q}$ is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $\boldsymbol{Q}(\Pi^*)$ to an estimand that is a function of the observed distributions.

Q = P*(y | do(x)) = ? Nope.

Proof. ∃ N_1, N_2

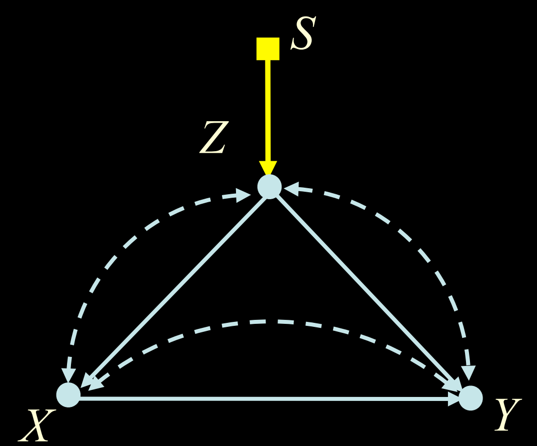

G

N_1

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow U_y \end{cases}$$

$P(u_i) = 1/2$

N_2

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow ((X \oplus Z \oplus U_{xy}) \wedge (Z \oplus U_{zy})) \vee U_y \end{cases}$$

$P(u_i) = 1/2$

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x)) = ?$ Nope.

Proof. $\exists\ N_1, N_2$ s.t. $G(N_1) = G(N_2) = G$,

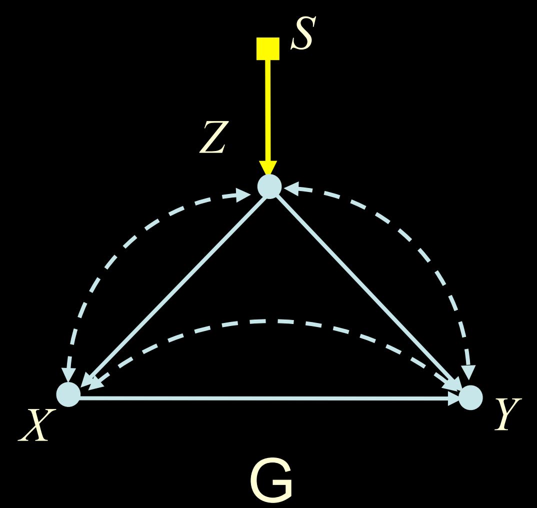

N_1

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow U_y \end{cases}$$

$P(u_i) = 1/2$

N_2

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow ((X \oplus Z \oplus U_{xy}) \wedge (Z \oplus U_{zy})) \vee U_y \end{cases}$$

$P(u_i) = 1/2$

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) if and only if there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x)) = ?$ Nope.

Proof. $\exists\ N_1, N_2$ s.t. $G(N_1) = G(N_2) = G$,
$P^*_1(z,x,y) = P^*_2(z,x,y)$, $P_1(z,x,y) = P_2(z,x,y)$,
$P_1(z,y|do(x)) = P_2(z,y|do(x))$, but

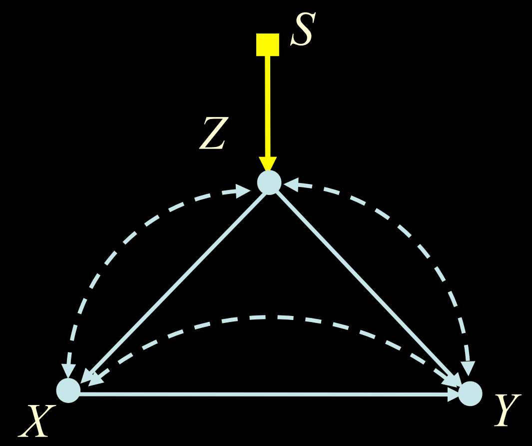

G

N1

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow U_y \end{cases}$$

$P(u_i) = 1/2$

N2

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow ((X \oplus Z \oplus U_{xy}) \wedge (Z \oplus U_{zy})) \vee U_y \end{cases}$$

$P(u_i) = 1/2$

RESULT 1: SYNTACTIC CHARACTERIZATION OF TRANSPORTABILITY

Thm: A causal quantity Q is transportable from Π to Π^* (G) **if and only if** there exists a do-calculus reduction of $Q(\Pi^*)$ to an estimand that is a function of the observed distributions.

$Q = P^*(y \mid do(x))$ = ? Nope.

Proof. $\exists$ N_1, N_2 s.t. $G(N_1) = G(N_2) = G$,
$P^*_1(z,x,y) = P^*_2(z,x,y)$, $P_1(z,x,y) = P_2(z,x,y)$,
$P_1(z,y|do(x)) = P_2(z,y|do(x))$, but
$P^*_1(y|do(x)) \neq P^*_2(y|do(x))$.

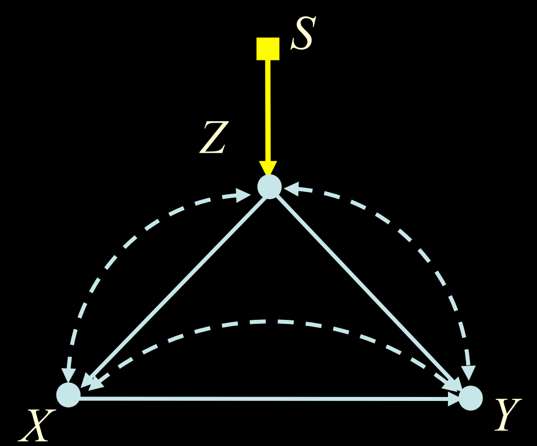

G

N_1

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow U_y \end{cases}$$

$P(u_i) = 1/2$

N_2

$$\begin{cases} Z \leftarrow U_{zy} \vee (S \wedge U_{xz}) \\ X \leftarrow Z \oplus U_{xy} \\ Y \leftarrow ((X \oplus Z \oplus U_{xy}) \wedge (Z \oplus U_{zy})) \vee U_y \end{cases}$$

$P(u_i) = 1/2$

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

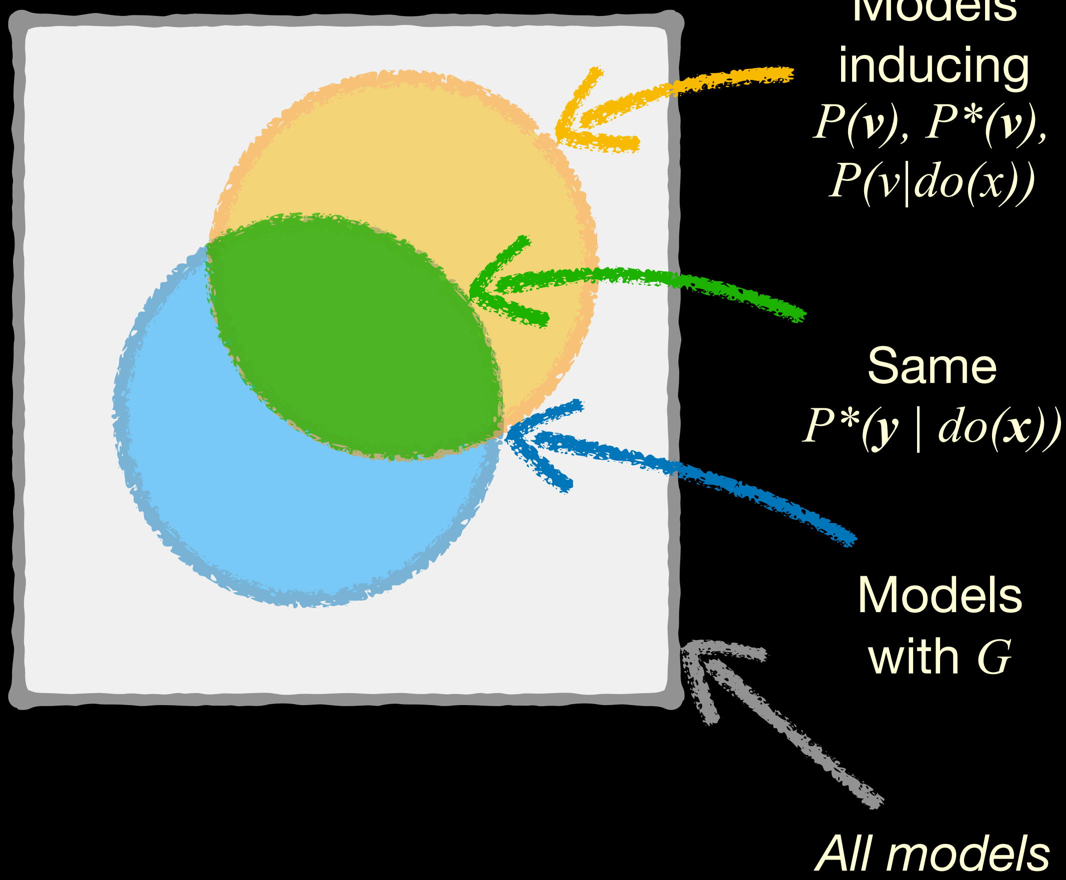

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

$P^*(y|do(x))$ is not transportable in G

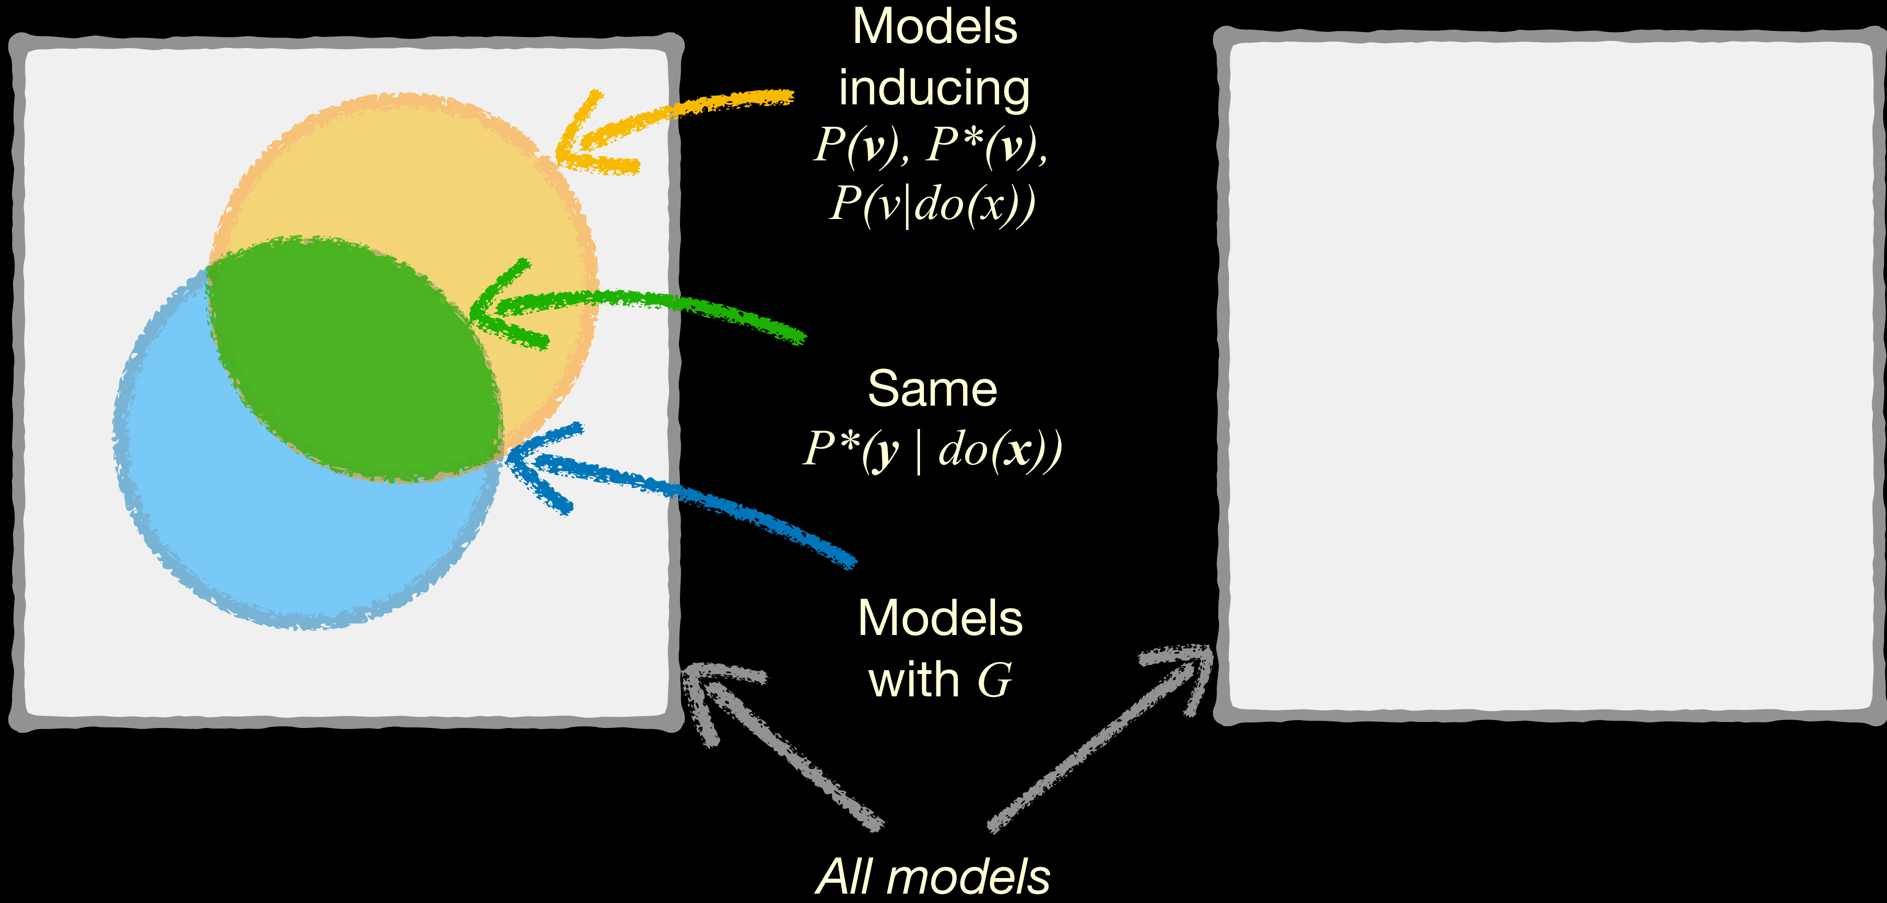

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

$P^*(y|do(x))$ is not transportable in G

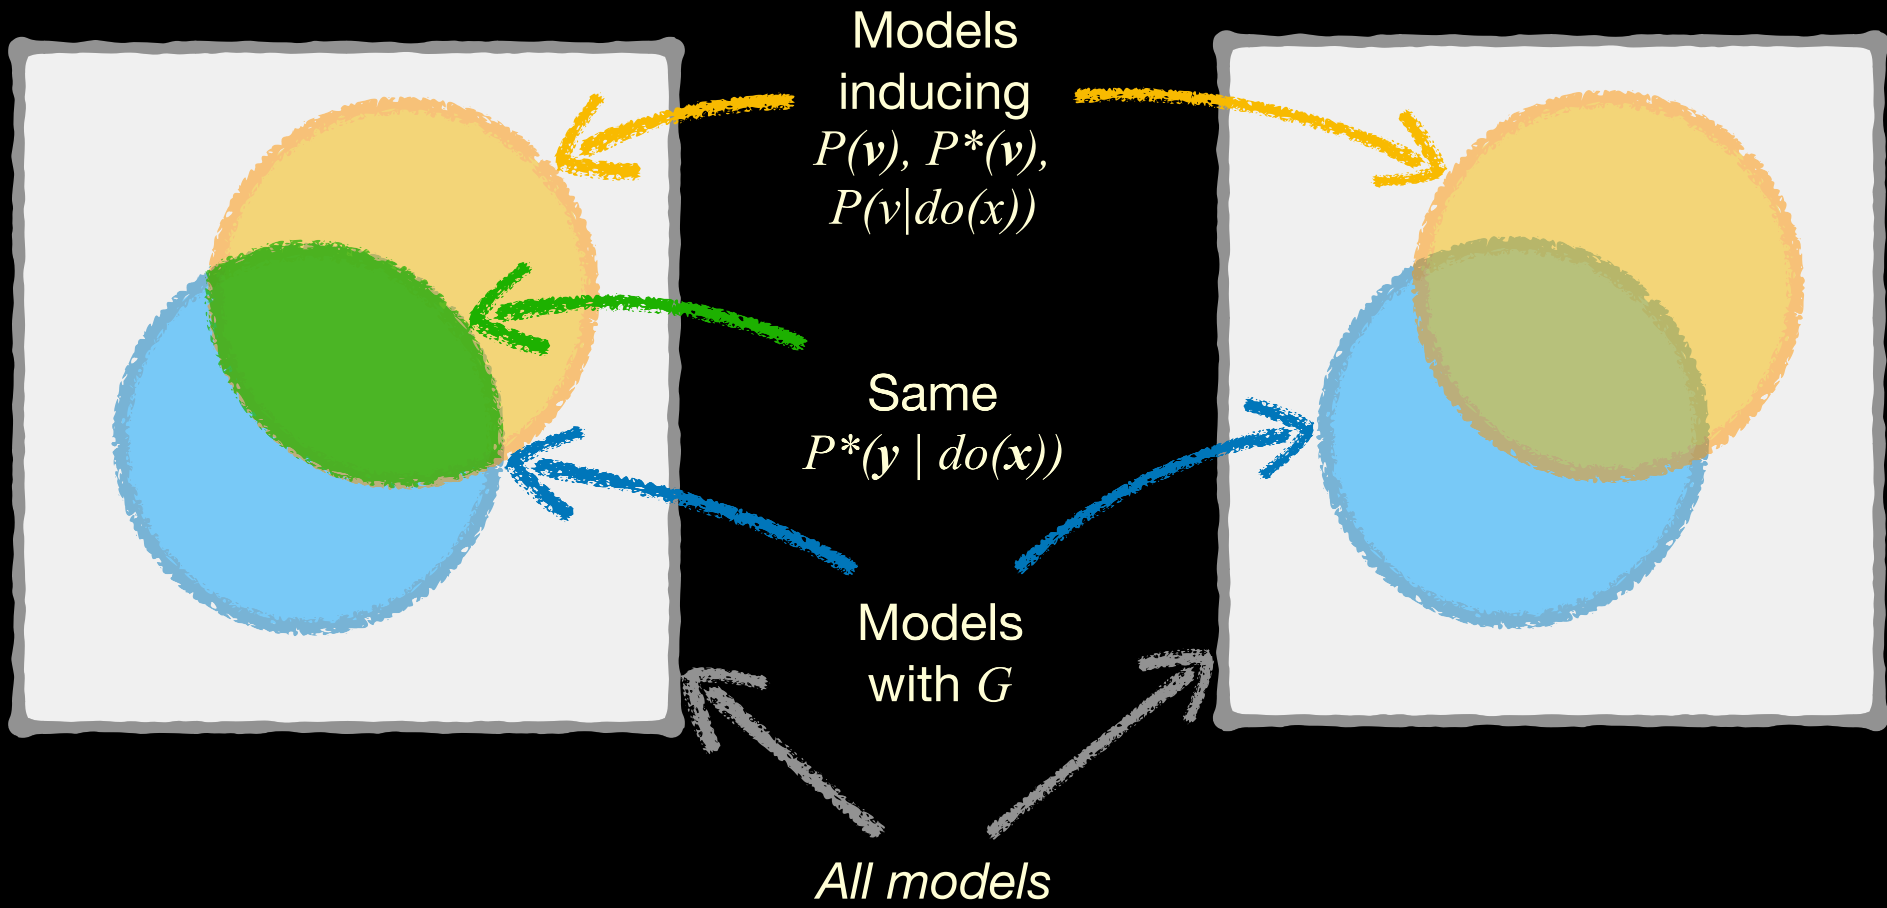

MEANING OF TRANSPORTABILITY

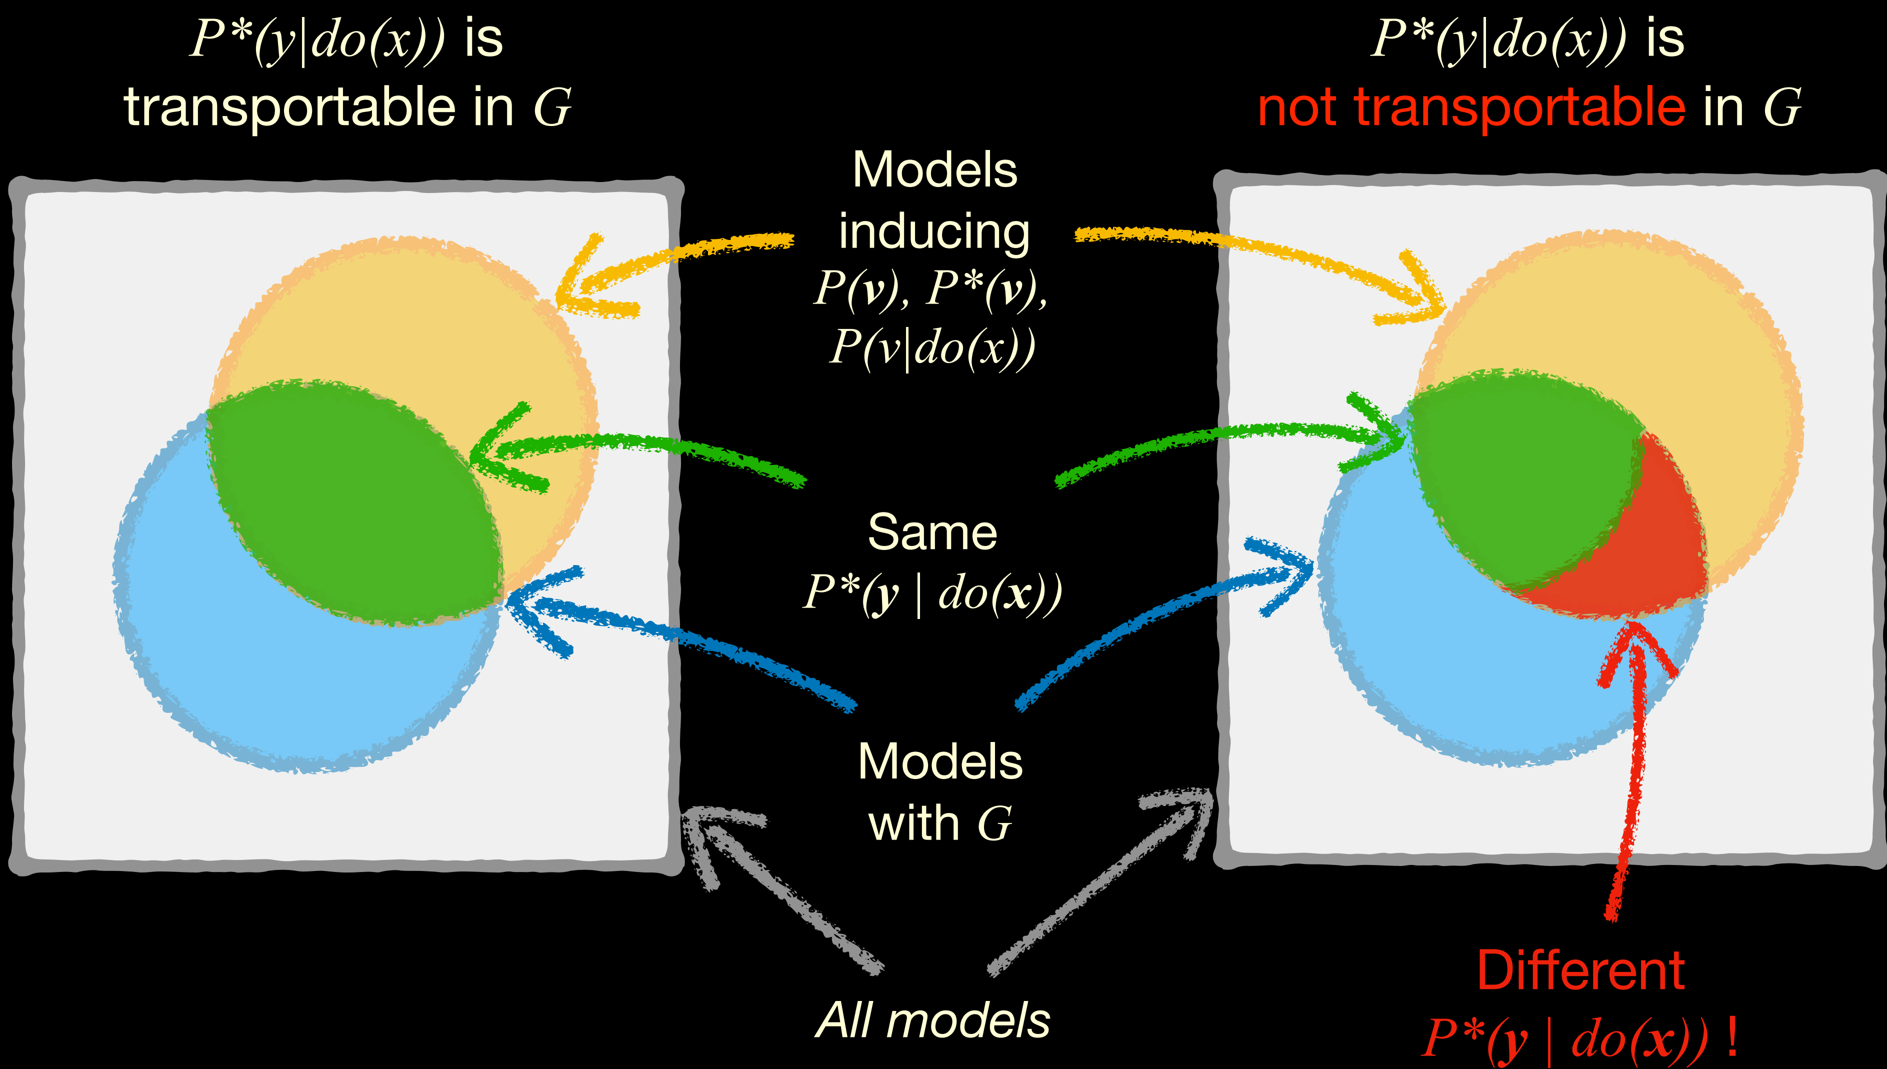

MEANING OF TRANSPORTABILITY

$P^*(y|do(x))$ is transportable in G

$P^*(y|do(x))$ is not transportable in G

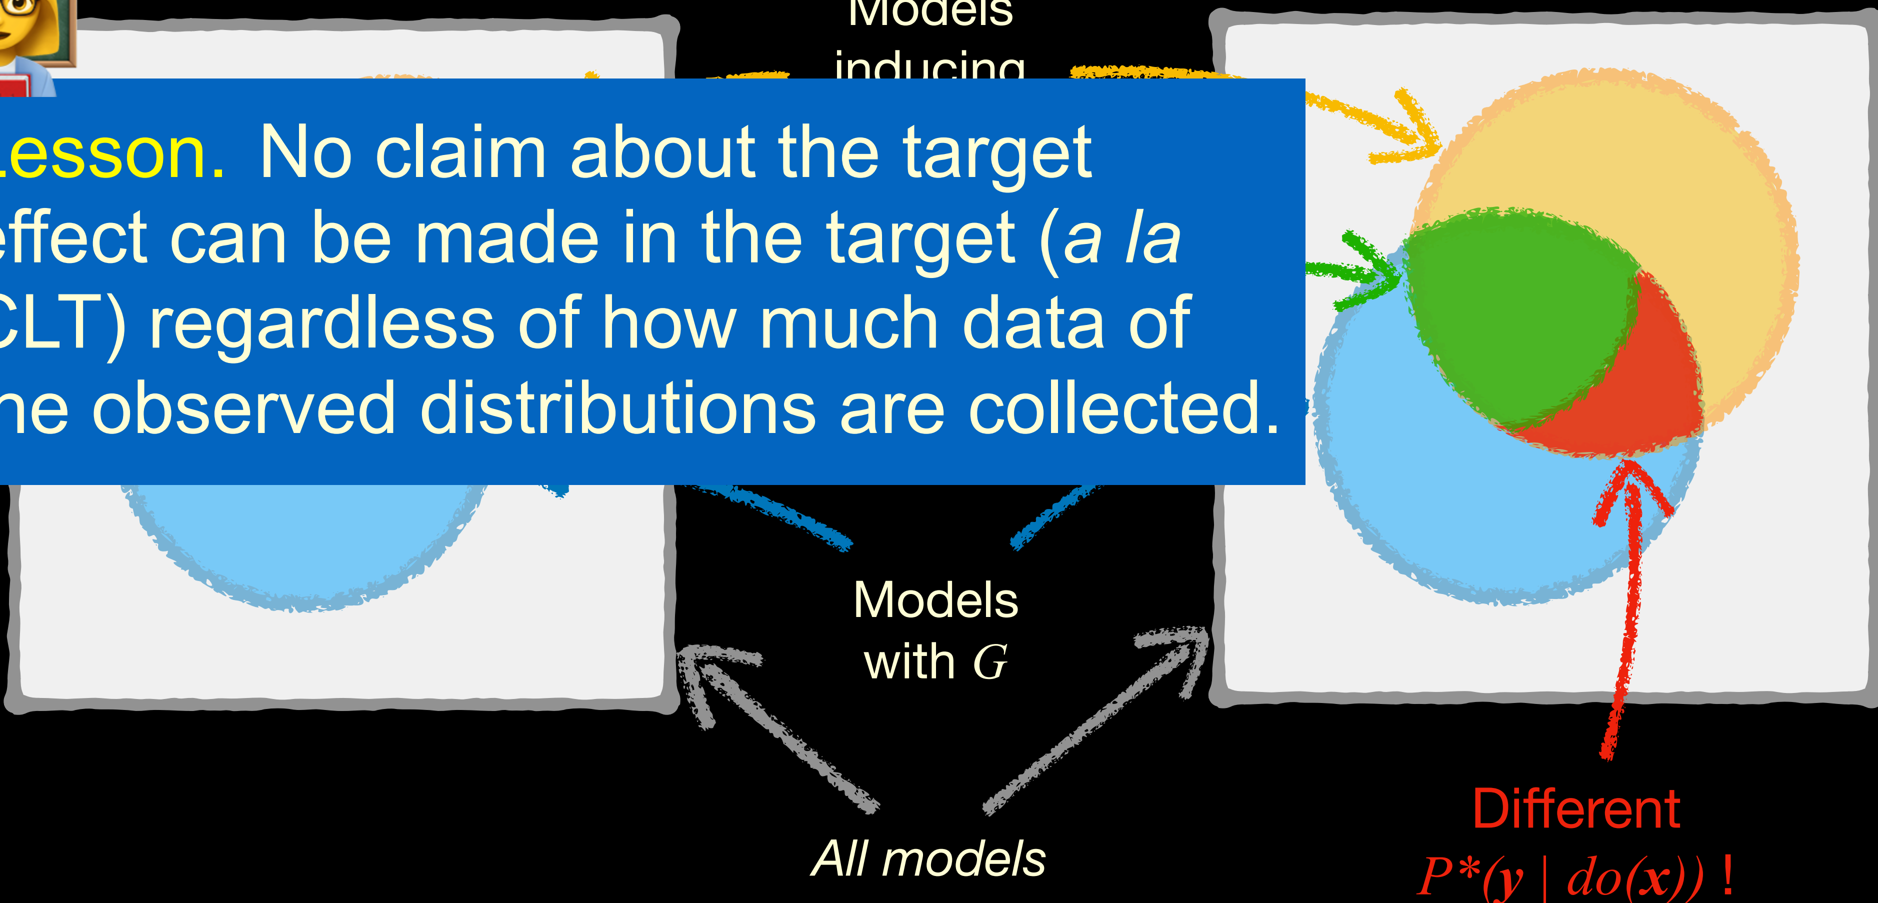

Lesson. No claim about the target effect can be made in the target (*a la* CLT) regardless of how much data of the observed distributions are collected.

RESULT 2: ALGORITHM TO DECIDE IF AN EFFECT IS TRANSPORTABLE

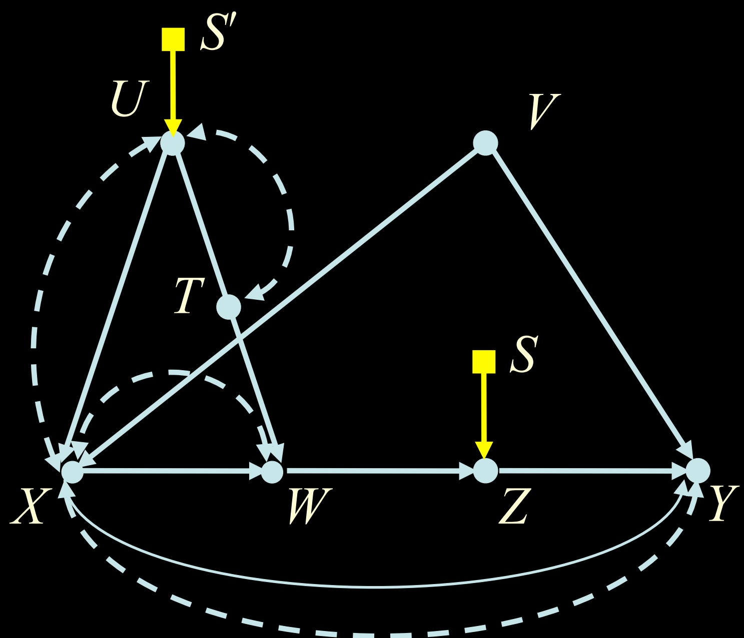

RESULT 2: ALGORITHM TO DECIDE IF AN EFFECT IS TRANSPORTABLE

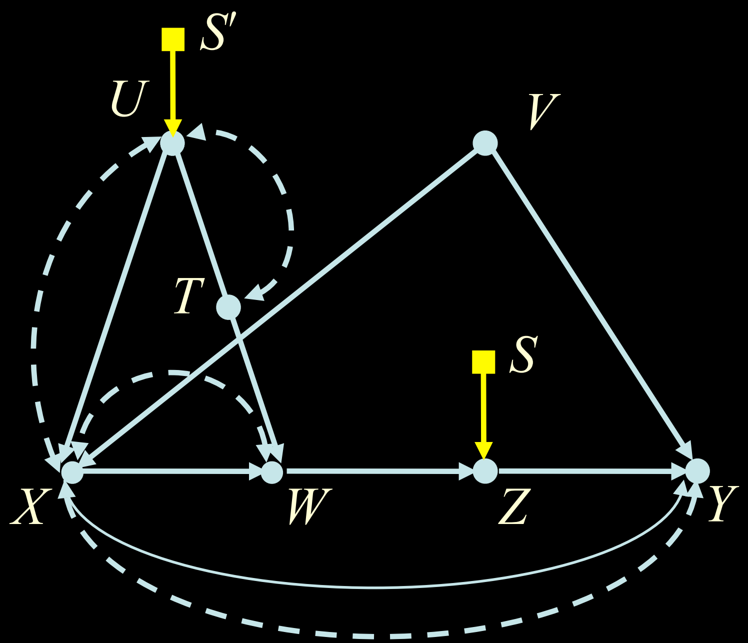

INPUT: Annotated Causal Graph

S Factors representing differences

OUTPUT:

1. Transportable or not?
2. Yes = Transport formula
    1. Measurements to be taken in the experimental domain
    2. Measurements to be taken in the target domain
3. No = Not possible

RESULT 2: ALGORITHM TO DECIDE IF AN EFFECT IS TRANSPORTABLE

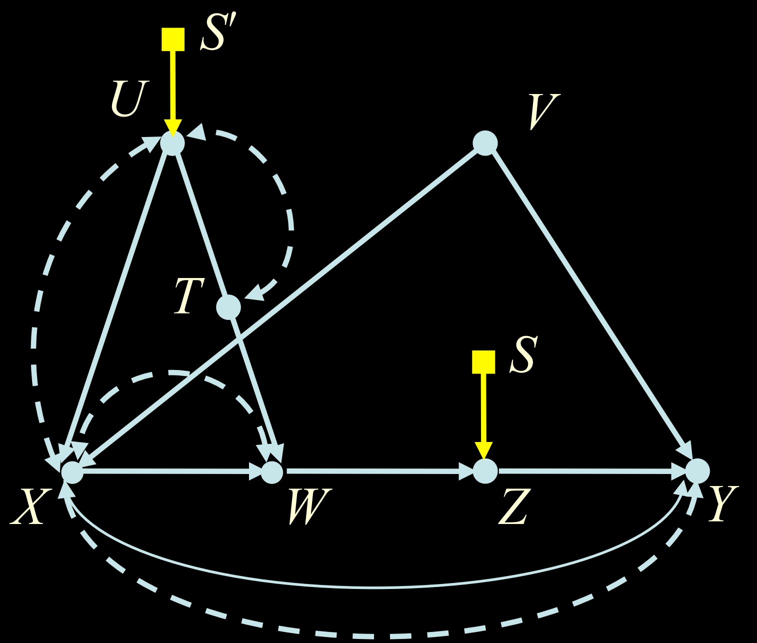

INPUT: Annotated Causal Graph

S → Factors representing differences

OUTPUT:

1. Transportable or not?
2. Yes = Transport formula
   1. Measurements to be taken in the experimental domain
   2. Measurements to be taken in the target domain
3. No = Not possible

$$P^*(y \mid do(x)) = \sum_z P(y \mid do(x), z) \sum_w P^*(z \mid w) \sum_t P(w \mid do(x), t)\, P^*(t)$$

RESULT 2: ALGORITHM TO DECIDE IF AN EFFECT IS TRANSPORTABLE

In reality…

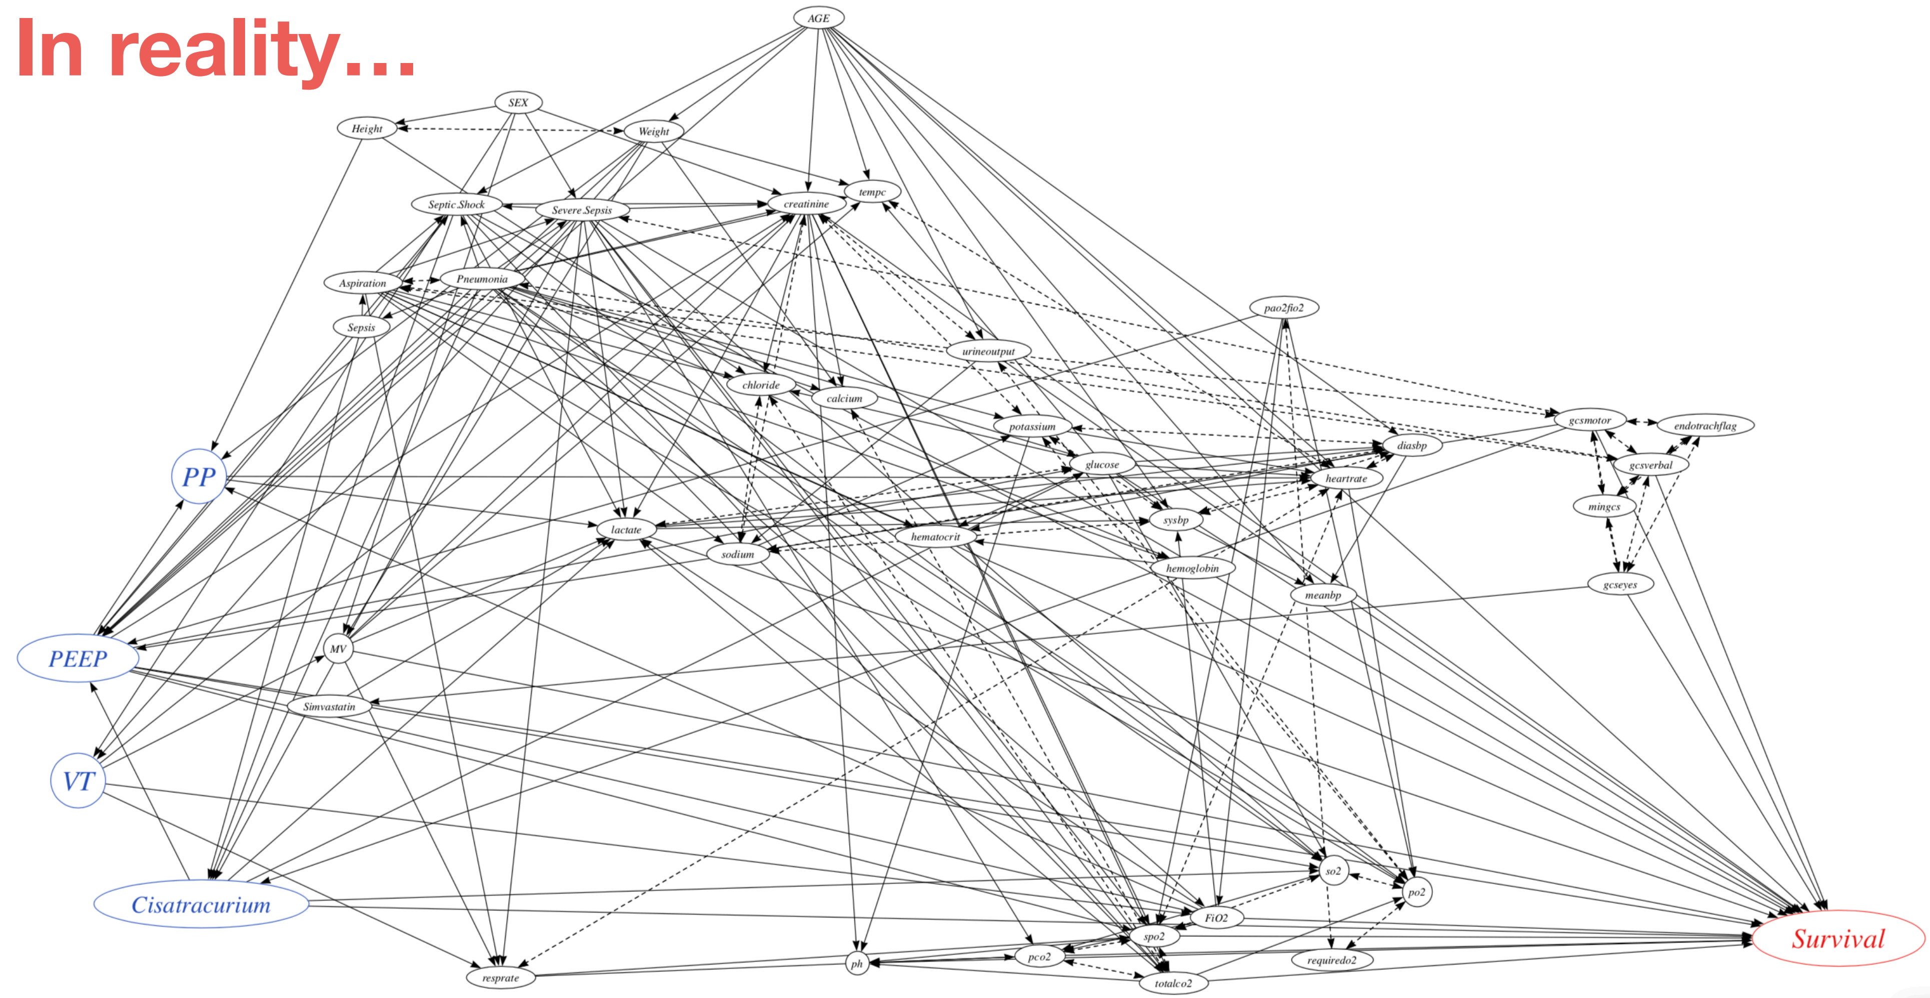

RESULT 2: ALGORITHM TO DECIDE IF AN EFFECT IS TRANSPO

Result:
- automated transportability analysis in large-scale settings.

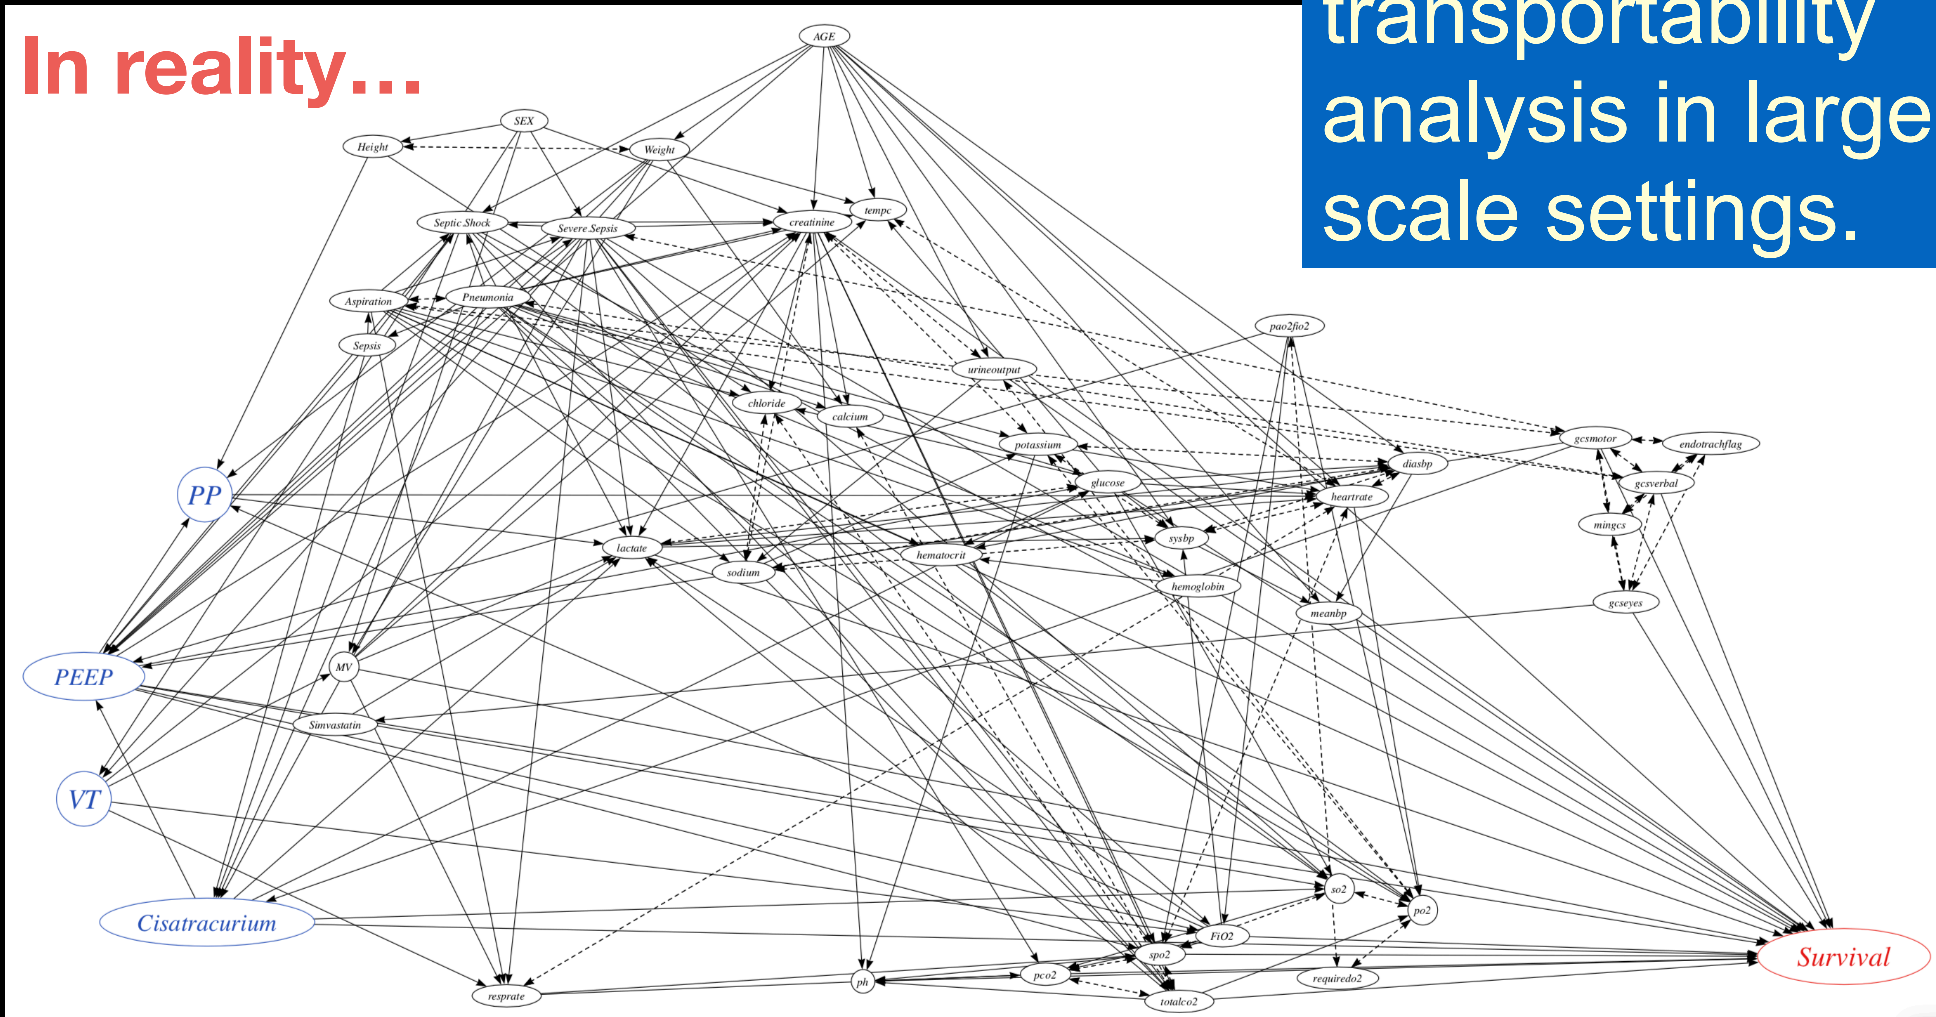

IS THE GOLD STANDARD GOLDEN?
(GENERALIZABILITY FROM CLINICAL TRIALS)

Before randomization

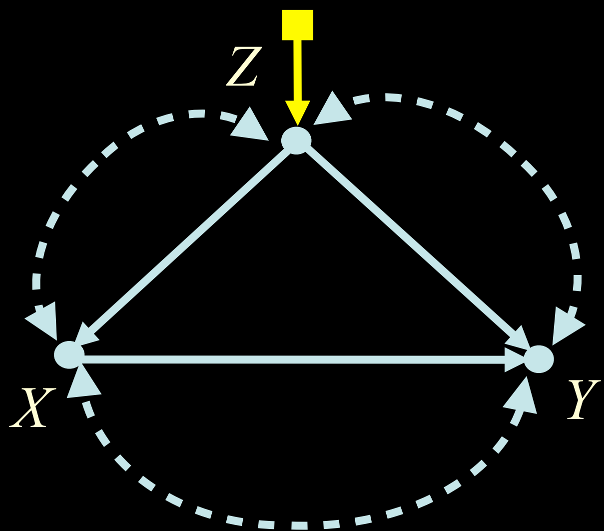

IS THE GOLD STANDARD GOLDEN?
(GENERALIZABILITY FROM CLINICAL TRIALS)

Before randomization

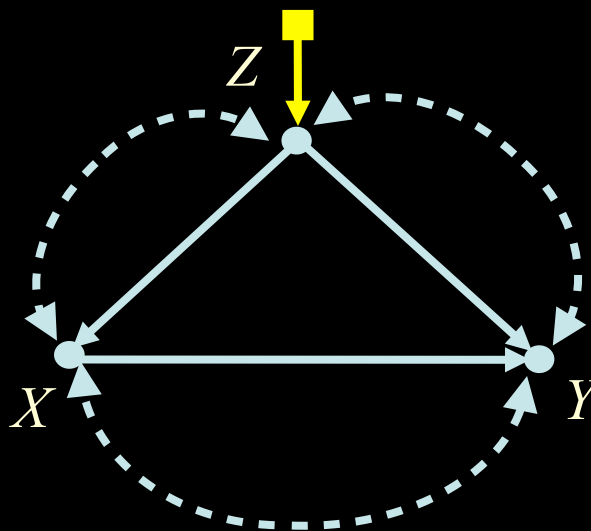

After randomization

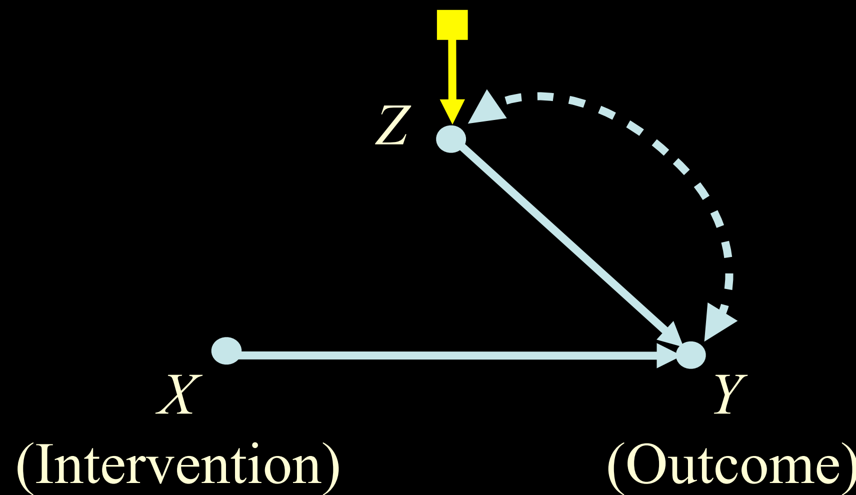

IS THE GOLD STANDARD GOLDEN?
(GENERALIZABILITY FROM CLINICAL TRIALS)

Before randomization

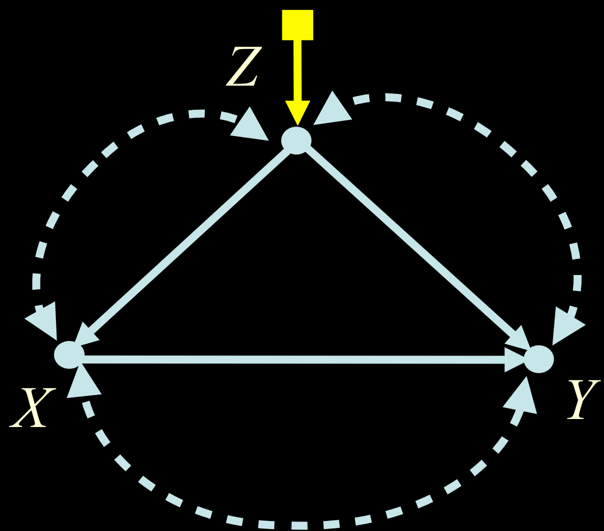

After randomization

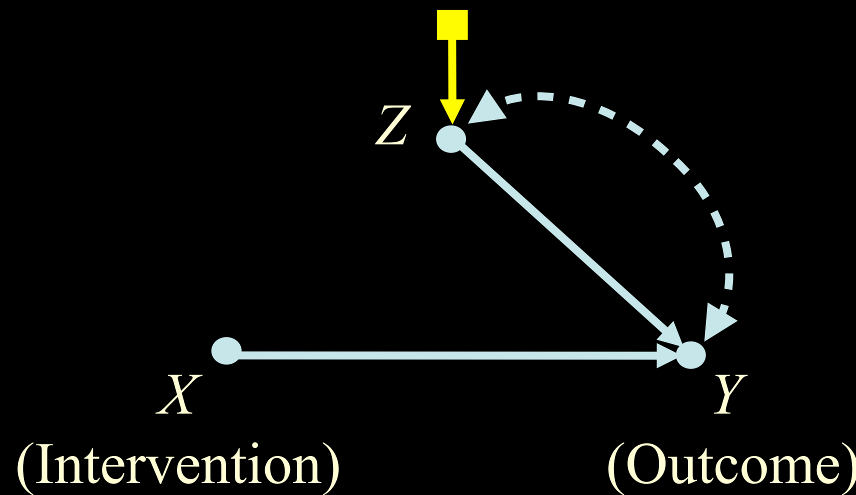

Lesson. Even if we have a perfect RCT, one still needs to go through a Transportability exercise. TR theory is unavoidable.

CausalAI LAB
Big Picture

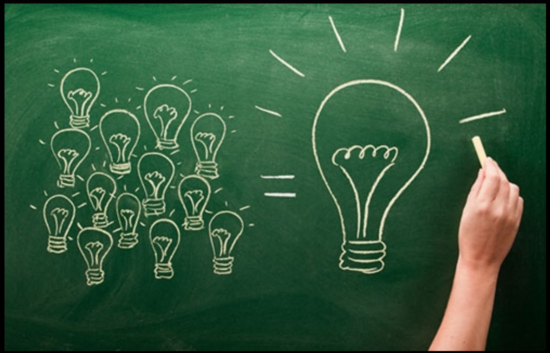

CausalAI Lab - Big Picture

CausalAI Lab - Big Picture

Postdoc Available

CausalAI Lab - Big Picture

Postdoc Available

Structural Causal Models

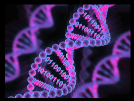

CausalAI Lab - Big Picture

Postdoc Available

Structural Causal Models

1. Explainability

(Effect identification and decomposition, Bias Analysis and Fairness, Robustness and Generalizability)

CausalAI Lab - Big Picture

Postdoc Available

Structural Causal Models

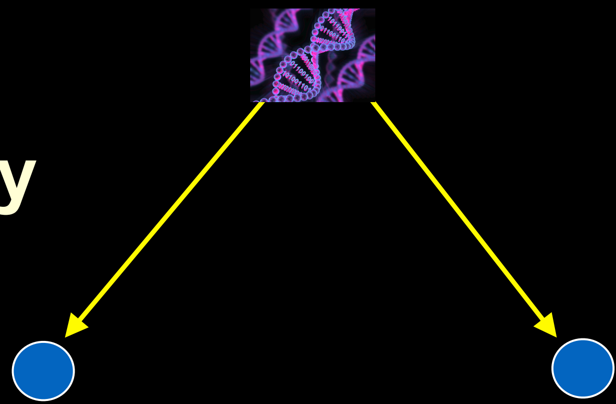

1. Explainability

(Effect identification and decomposition, Bias Analysis and Fairness, Robustness and Generalizability)

2. Decision-Making

(Reinforcement Learning, Randomized Clinical Trials, Personalized Decision-Making)

CausalAI Lab - Big Picture

Postdoc Available

Structural Causal Models

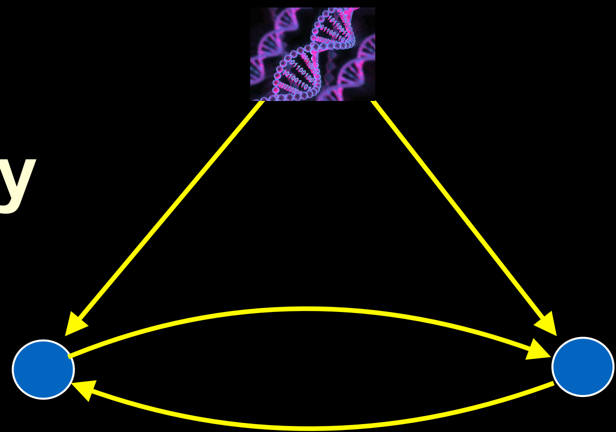

1. Explainability

(Effect identification and decomposition, Bias Analysis and Fairness, Robustness and Generalizability)

2. Decision-Making

(Reinforcement Learning, Randomized Clinical Trials, Personalized Decision-Making)

CausalAI Lab - Big Picture

Postdoc Available

Structural Causal Models

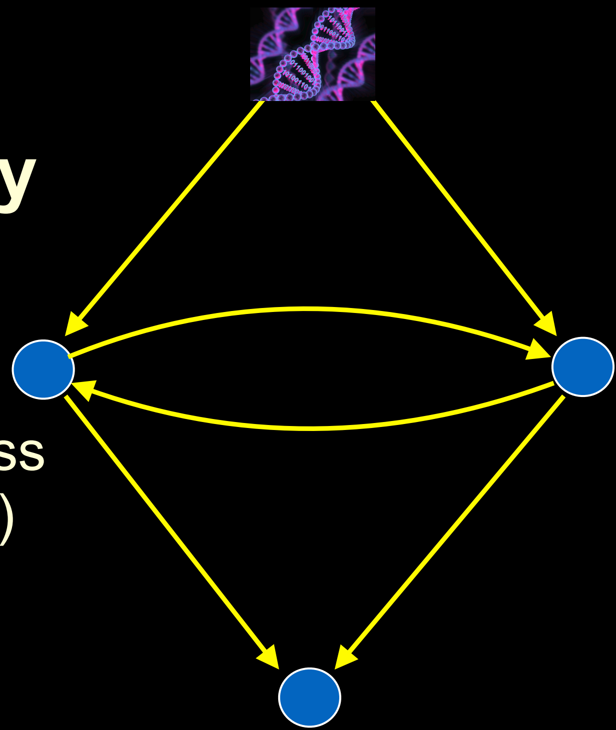

1. Explainability

(Effect identification and decomposition, Bias Analysis and Fairness, Robustness and Generalizability)

2. Decision-Making

(Reinforcement Learning, Randomized Clinical Trials, Personalized Decision-Making)

3. Applications, Education, Software

CausalAI Lab - Big Picture

Postdoc Available

Structural Causal Models

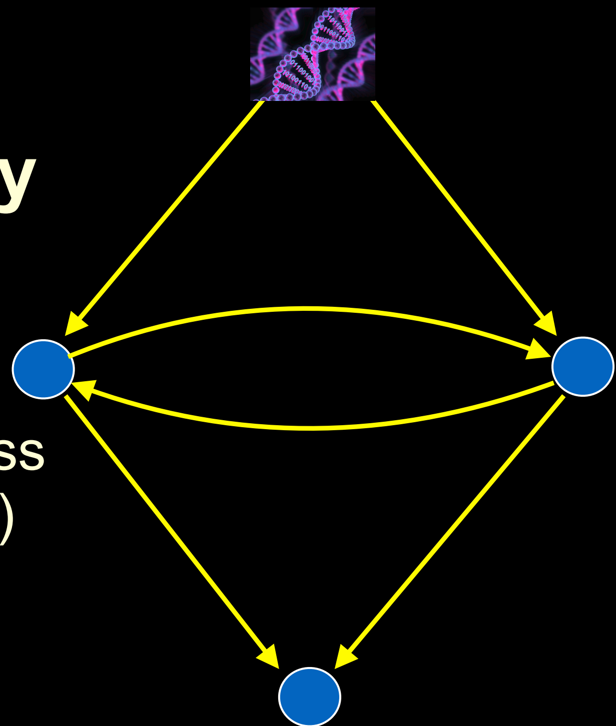

1. Explainability

(Effect identification and decomposition, Bias Analysis and Fairness, Robustness and Generalizability)

2. Decision-Making

(Reinforcement Learning, Randomized Clinical Trials, Personalized Decision-Making)

3. Applications, Education, Software

AI-ML:

Principles and tools for designing robust and adaptable learning systems.

CausalAI Lab - Big Picture

Postdoc Available

Structural Causal Models

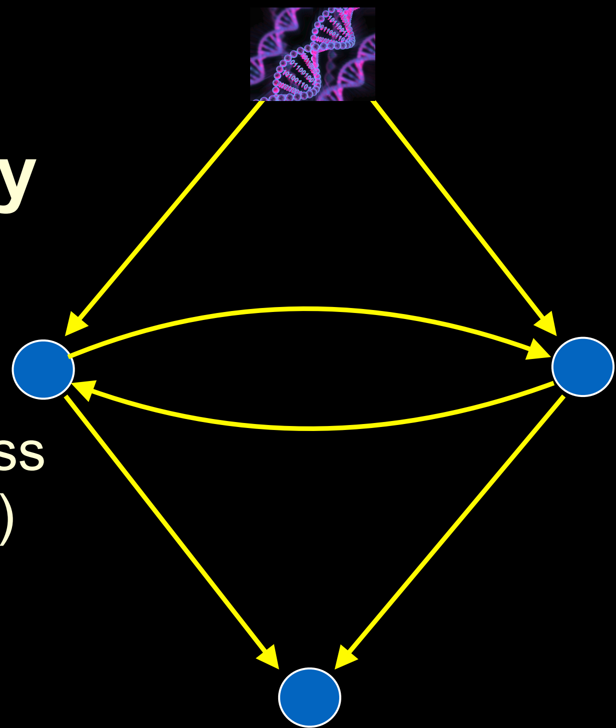

1. Explainability

(Effect identification and decomposition, Bias Analysis and Fairness, Robustness and Generalizability)

2. Decision-Making

(Reinforcement Learning, Randomized Clinical Trials, Personalized Decision-Making)

3. Applications, Education, Software

Data Science:
Principled ("scientific") inferences from large data collections.

AI-ML:
Principles and tools for designing robust and adaptable learning systems.

CausalAI Lab - Big Picture

Postdoc Available

Structural Causal Models

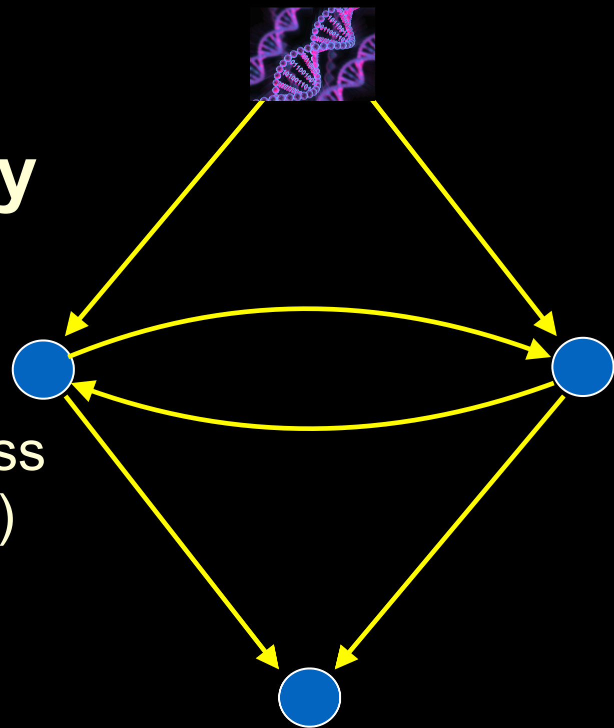

1. Explainability

(Effect identification and decomposition, Bias Analysis and Fairness, Robustness and Generalizability)

2. Decision-Making

(Reinforcement Learning, Randomized Clinical Trials, Personalized Decision-Making)

3. Applications, Education, Software

Data Science:
Principled ("scientific") inferences from large data collections.

AI-ML:
Principles and tools for designing robust and adaptable learning systems.

Other recent results in CausalAI
I haven't talked about today…

Other recent results in CausalAI I haven’t talked about today…

- Where to intervene (NeurIPS’18, AAAI’19) — Given a causal model, how to design an intervention that brings about the desired change?

Other recent results in CausalAI I haven't talked about today…

- Where to intervene (NeurIPS'18, AAAI'19) — Given a causal model, how to design an intervention that brings about the desired change?
- Approximate Causal Inference (UAI'18, IJCAI'18) — How to perform causal analysis when the model is too weak or just an equivalence class of causal models is available?

Other recent results in CausalAI I haven't talked about today…

- Where to intervene (NeurIPS'18, AAAI'19) — Given a causal model, how to design an intervention that brings about the desired change?
- Approximate Causal Inference (UAI'18, IJCAI'18) — How to perform causal analysis when the model is too weak or just an equivalence class of causal models is available?
- Causal Fairness Analysis (NeurIPS'18, AAAI'18) — How to determine (and then quantify, correct for) whether a certain decision-making process is unfair when the causal mechanisms are unknown?

Other recent results in CausalAI I haven't talked about today…

- Where to intervene (NeurIPS'18, AAAI'19) — Given a causal model, how to design an intervention that brings about the desired change?
- Approximate Causal Inference (UAI'18, IJCAI'18) — How to perform causal analysis when the model is too weak or just an equivalence class of causal models is available?
- Causal Fairness Analysis (NeurIPS'18, AAAI'18) — How to determine (and then quantify, correct for) whether a certain decision-making process is unfair when the causal mechanisms are unknown?
- Model learning from heterogenous sources (NeurIPS'17, ICML'18) — How to construct a causal model from a combination of heterogenous observational and experimental data?
- Counterfactual decision-making (NeurIPS'15, ICML'17, AAAI'19) — How to make individual-level decisions from population-level data?

CONCLUSIONS

CONCLUSIONS

- **Causal data science:** principled way of combining substantive knowledge about the phenomenon under investigation with data to generate causal explanations and better decision-making.

CONCLUSIONS

- Causal data science: principled way of combining substantive knowledge about the phenomenon under investigation with data to generate causal explanations and better decision-making.

CONCLUSIONS

- Causal data science: principled way of combining substantive knowledge about the phenomenon under investigation with data to generate causal explanations and better decision-making.
- "Not all data is created equal" - Most modern data collections are messy and plagued with systematic biases that need to be understood and controlled so that causal claims can be made. Causal data science solves long-standing challenges in the empirical sciences, including confounding bias, selection bias, generalizability of experimental studies, and data-fusion.

CONCLUSIONS

- Causal data science: principled way of combining substantive knowledge about the phenomenon under investigation with data to generate causal explanations and better decision-making.
- "Not all data is created equal" - Most modern data collections are messy and plagued with systematic biases that need to be understood and controlled so that causal claims can be made. Causal data science solves long-standing challenges in the empirical sciences, including confounding bias, selection bias, generalizability of experimental studies, and data-fusion.
- Future → Automated scientist

CONCLUSIONS

- Causal data science: principled way of combining substantive knowledge about the phenomenon under investigation with data to generate causal explanations and better decision-making.
- “Not all data is created equal” - Most modern data collections are messy and plagued with systematic biases that need to be understood and controlled so that causal claims can be made. Causal data science solves long-standing challenges in the empirical sciences, including confounding bias, selection bias, generalizability of experimental studies, and data-fusion.
- Future → Automated scientist
  - Reasoning with cause and effect relationships is part of the core of the ‘scientific method’ and has been understood in great generality;

CONCLUSIONS

- Causal data science: principled way of combining substantive knowledge about the phenomenon under investigation with data to generate causal explanations and better decision-making.

- “Not all data is created equal” - Most modern data collections are messy and plagued with systematic biases that need to be understood and controlled so that causal claims can be made. Causal data science solves long-standing challenges in the empirical sciences, including confounding bias, selection bias, generalizability of experimental studies, and data-fusion.

- Future → Automated scientist
  - Reasoning with cause and effect relationships is part of the core of the ‘scientific method’ and has been understood in great generality;

CONCLUSIONS

- Causal data science: principled way of combining substantive knowledge about the phenomenon under investigation with data to generate causal explanations and better decision-making.
- "Not all data is created equal" - Most modern data collections are messy and plagued with systematic biases that need to be understood and controlled so that causal claims can be made. Causal data science solves long-standing challenges in the empirical sciences, including confounding bias, selection bias, generalizability of experimental studies, and data-fusion.
- Future → Automated scientist
  - Reasoning with cause and effect relationships is part of the core of the 'scientific method' and has been understood in great generality;
  - Strategy: combine causal inference theory with AI-ML techniques to automate the process and close the scientific discovery loop.

"Development of Western Science is based on two great achievements, the invention of the formal logical system (in Euclidean geometry) by the Greek philosophers, and the discovery of the possibility to find out causal relationships by systematic experiment (during the Renaissance)". Albert Einstein.

"Imagine how much harder physics would be if electrons had feelings!". Richard Feynman.

THANK YOU!

CAUSAL FAIRNESS ANALYSIS

NeurIPS'18
AAAI'19

CAUSAL FAIRNESS ANALYSIS

- Context. AI-ML systems are starting to replace humans in critical decision-making infra-structure, including bank loans, hiring decisions, and criminal incarceration.

CAUSAL FAIRNESS ANALYSIS

- Context. AI-ML systems are starting to replace humans in critical decision-making infra-structure, including bank loans, hiring decisions, and criminal incarceration.
- It's natural to envision a future where most of society's decision-making infra-structure is delegated to automated systems.

CAUSAL FAIRNESS ANALYSIS

- Context. AI-ML systems are starting to replace humans in critical decision-making infra-structure, including bank loans, hiring decisions, and criminal incarceration.
- It's natural to envision a future where most of society's decision-making infra-structure is delegated to automated systems.
- Consider an example in the justice system.

CAUSAL FAIRNESS ANALYSIS

- Context. AI-ML systems are starting to replace humans in critical decision-making infra-structure, including bank loans, hiring decisions, and criminal incarceration.
- It's natural to envision a future where most of society's decision-making infra-structure is delegated to automated systems.
- Consider an example in the justice system.

Judges

CAUSAL FAIRNESS ANALYSIS

- Context. AI-ML systems are starting to replace humans in critical decision-making infra-structure, including bank loans, hiring decisions, and criminal incarceration.
- It's natural to envision a future where most of society's decision-making infra-structure is delegated to automated systems.
- Consider an example in the justice system.

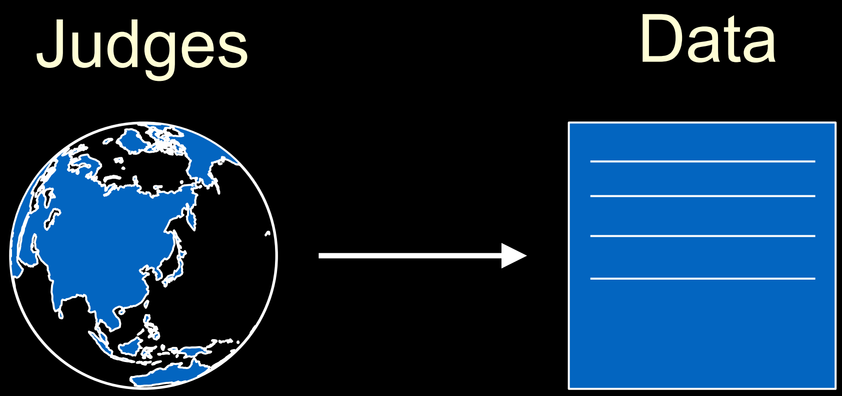

NeurIPS’18
AAAI’19

CAUSAL FAIRNESS ANALYSIS

- Context. AI-ML systems are starting to replace humans in critical decision-making infra-structure, including bank loans, hiring decisions, and criminal incarceration.
- It’s natural to envision a future where most of society’s decision-making infra-structure is delegated to automated systems.
- Consider an example in the justice system.

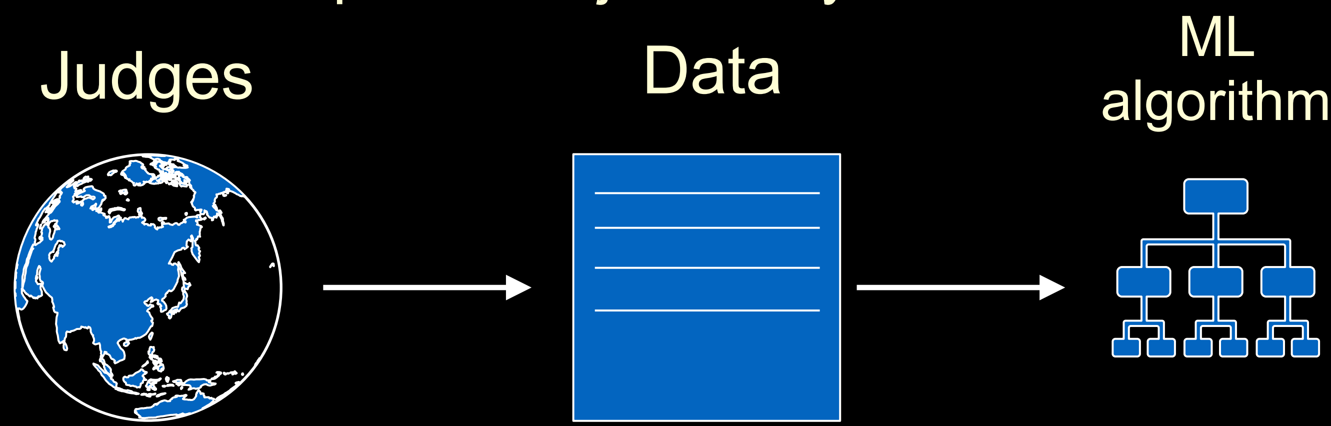

CAUSAL FAIRNESS ANALYSIS

NeurIPS'18
AAAI'19

- Context. AI-ML systems are starting to replace humans in critical decision-making infra-structure, including bank loans, hiring decisions, and criminal incarceration.
- It's natural to envision a future where most of society's decision-making infra-structure is delegated to automated systems.
- Consider an example in the justice system.

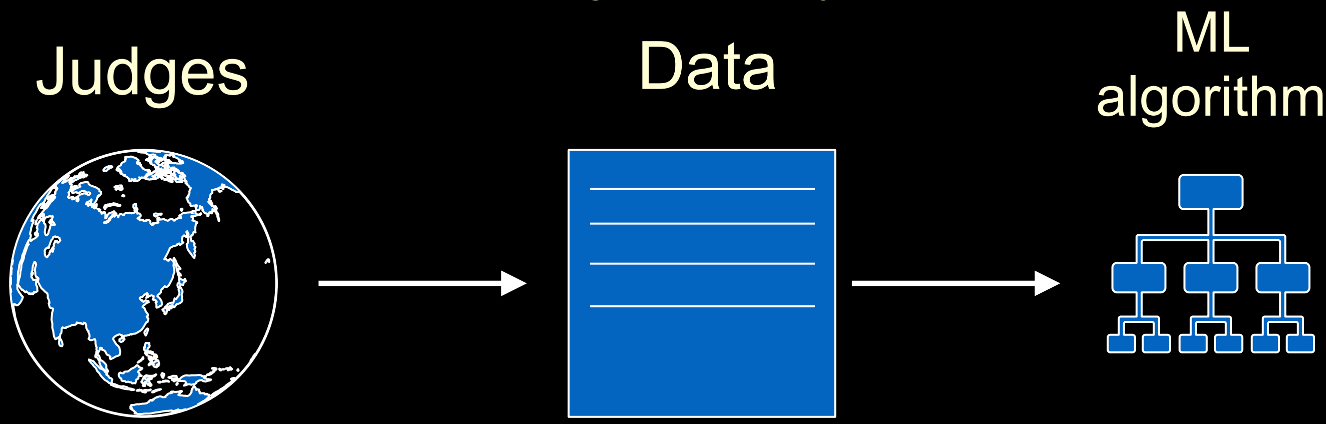

- (Possibly) surprising observations:

CAUSAL FAIRNESS ANALYSIS

- Context. AI-ML systems are starting to replace humans in critical decision-making infra-structure, including bank loans, hiring decisions, and criminal incarceration.
- It's natural to envision a future where most of society's decision-making infra-structure is delegated to automated systems.
- Consider an example in the justice system.

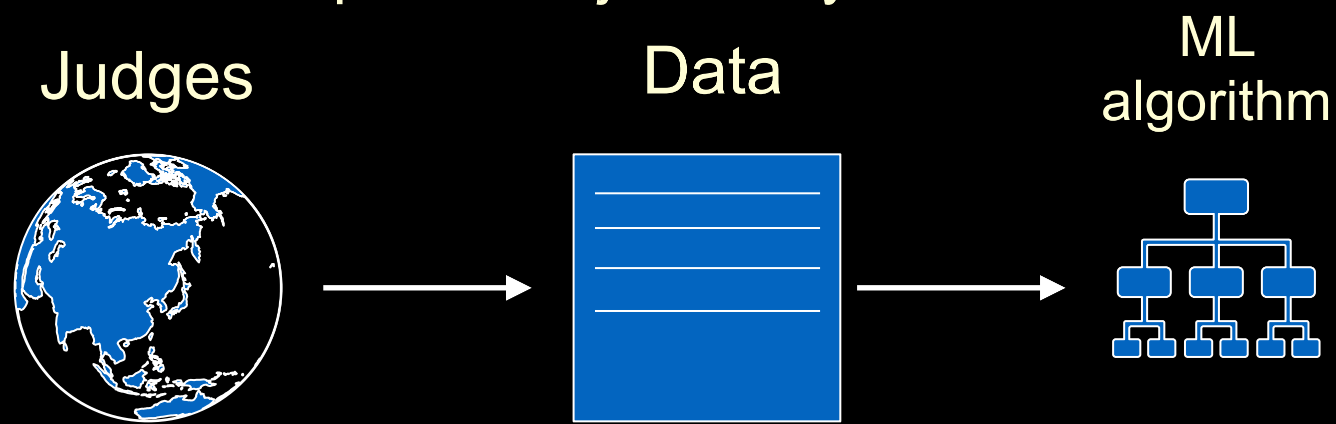

- (Possibly) surprising observations:
  1. The data reflects the judges decisions, which includes their biases.
  2. The trained ML algorithm tries to "mimic" the patterns in the data.

CAUSAL FAIRNESS ANALYSIS

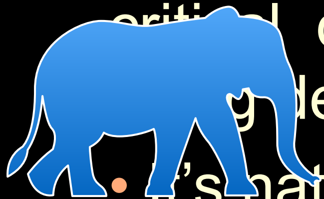

- Context. AI-ML systems are starting to replace humans in
decision-making infra-structure, including bank loans
g dec

- It's natur
making in

- Conside

Question. If the data cannot be used as baseline (or oracle), what can we, machine learners, really do?

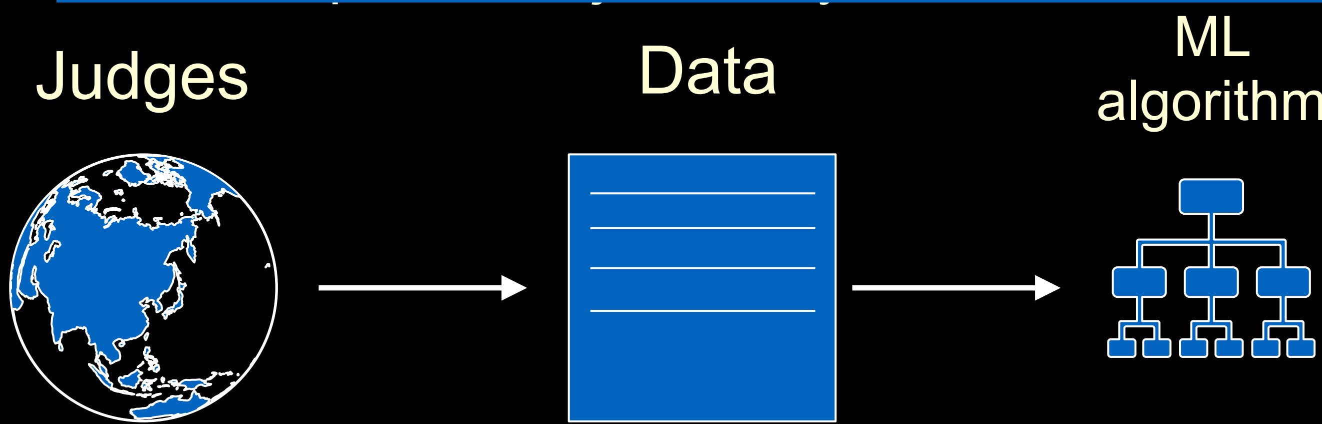

- (Possibly) surprising observations:
  1. The data reflects the judges decisions, which includes their biases.
  2. The trained ML algorithm tries to "mimic" the patterns in the data.

CAUSAL FAIRNESS ANALYSIS

CAUSAL FAIRNESS ANALYSIS

- Law (e.g., Title VII): What makes some action "discriminatory" are the "reasons" that lead to the specific disparity, not the disparity itself.

CAUSAL FAIRNESS ANALYSIS

- Law (e.g., Title VII): What makes some action "discriminatory" are the "reasons" that lead to the specific disparity, not the disparity itself.
    - The plaintiff must demonstrate "a strong causal connection" between the alleged discriminatory practice and the observed statistical disparity, otherwise the case will be dismissed (Texas Dept. of Housing and Community Affairs v. Inclusive Communities Project, Inc., 576 U.S. __ (2015)).

CAUSAL FAIRNESS ANALYSIS

- Law (e.g., Title VII): What makes some action "discriminatory" are the "reasons" that lead to the specific disparity, not the disparity itself.
  - The plaintiff must demonstrate "a strong causal connection" between the alleged discriminatory practice and the observed statistical disparity, otherwise the case will be dismissed (Texas Dept. of Housing and Community Affairs v. Inclusive Communities Project, Inc., 576 U.S. __ (2015)).

CAUSAL FAIRNESS ANALYSIS

- Law (e.g., Title VII): What makes some action "discriminatory" are the "reasons" that lead to the specific disparity, not the disparity itself.
    - The plaintiff must demonstrate "a strong causal connection" between the alleged discriminatory practice and the observed statistical disparity, otherwise the case will be dismissed (Texas Dept. of Housing and Community Affairs v. Inclusive Communities Project, Inc., 576 U.S. __ (2015)).

CAUSAL FAIRNESS ANALYSIS

- Law (e.g., Title VII): What makes some action "discriminatory" are the "reasons" that lead to the specific disparity, not the disparity itself.
  - The plaintiff must demonstrate "a strong causal connection" between the alleged discriminatory practice and the observed statistical disparity, otherwise the case will be dismissed (Texas Dept. of Housing and Community Affairs v. Inclusive Communities Project, Inc., 576 U.S. __ (2015)).

- Question. How can we disentangle "the reasons" for the judges to be acting in the way they did from the observed reality and data? How to measure the underlying causal mechanisms that are unobserved?

CAUSAL FAIRNESS ANALYSIS

- Example: Assume X is the protected attribute, e.g., religious beliefs, Y is the outcome of interest, hiring decision, and Z and W are covariates encoding, respectively, the educational level and proximity to work of the applicant.

Causal Model

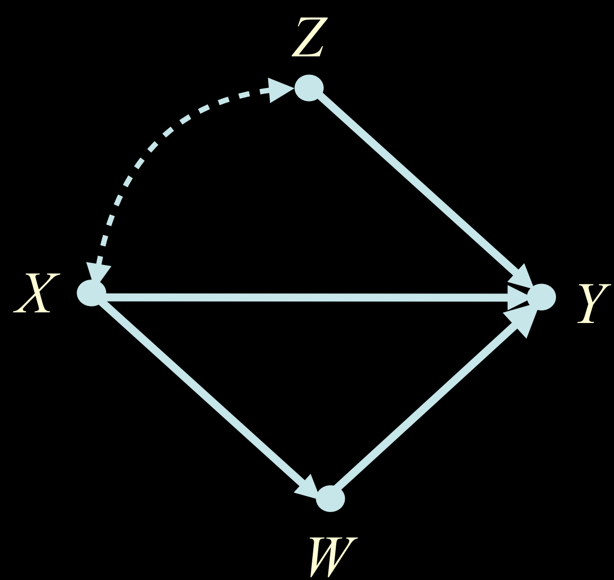

CAUSAL FAIRNESS ANALYSIS

- Example: Assume X is the protected attribute, e.g., religious beliefs, Y is the outcome of interest, hiring decision, and Z and W are covariates encoding, respectively, the educational level and proximity to work of the applicant.

Causal Model

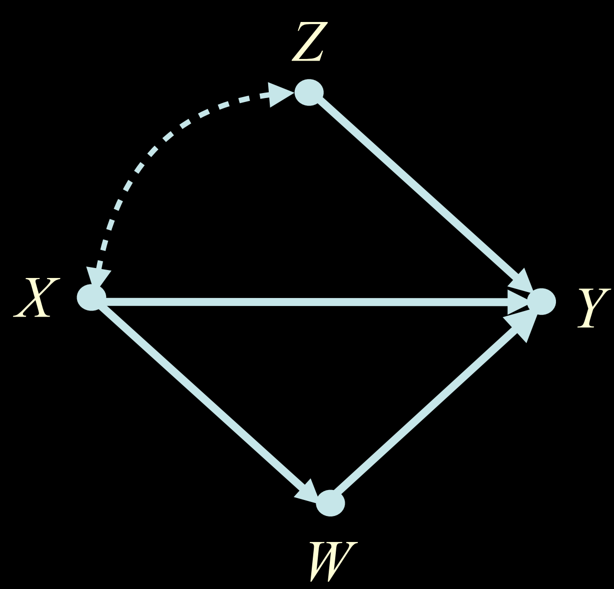

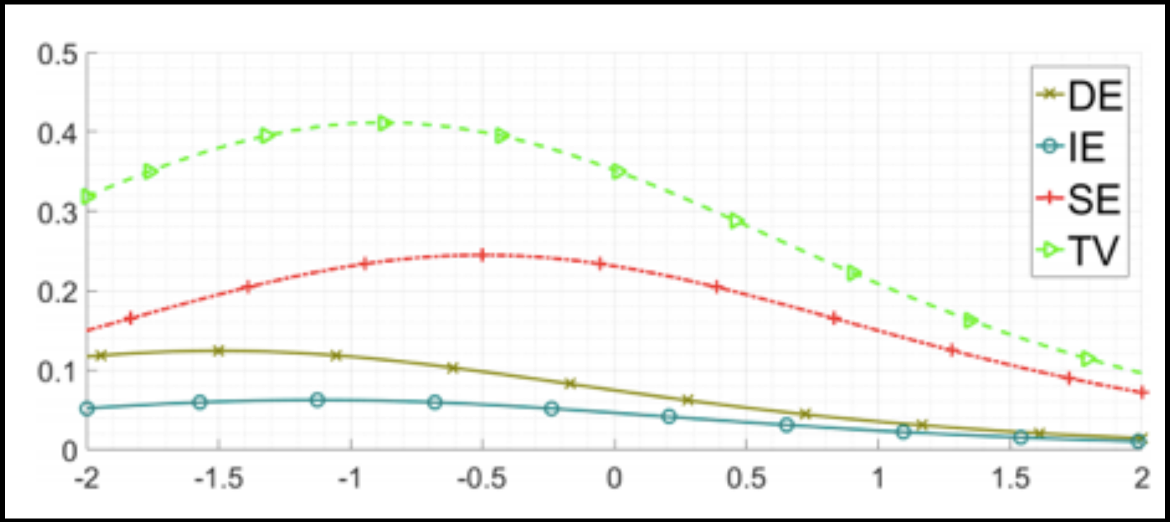

CAUSAL FAIRNESS ANALYSIS

NeurIPS'18
AAAI'19

- **Example:** Assume X is the protected attribute, e.g., religious beliefs, Y is the outcome of interest, hiring decision, and Z and W are covariates encoding, respectively, the educational level and proximity to work of the applicant.

Causal Model

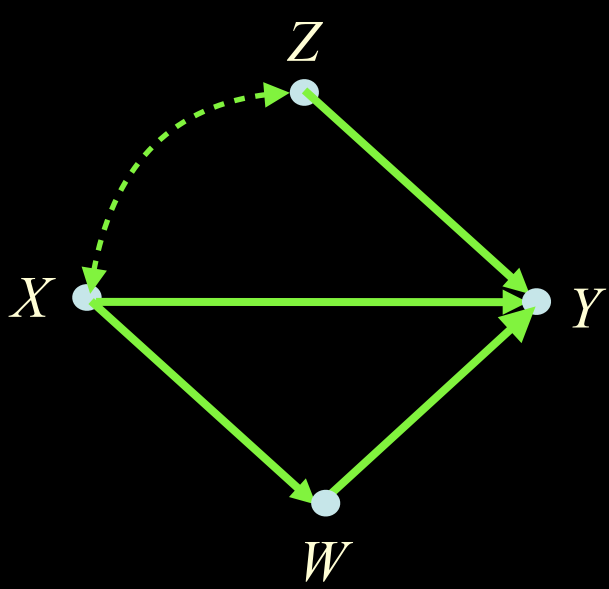

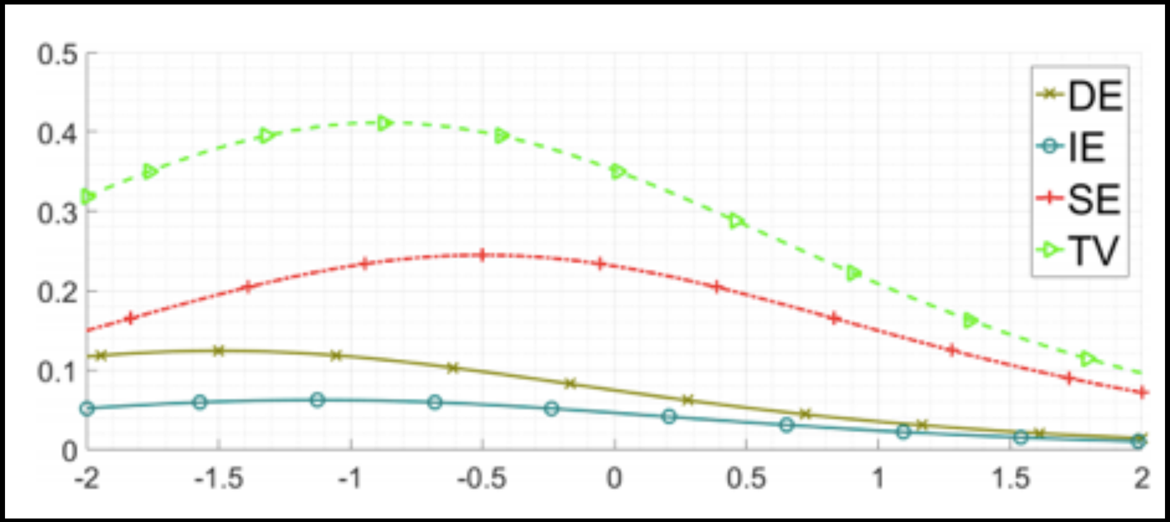

CAUSAL FAIRNESS ANALYSIS

- **Example:** Assume X is the protected attribute, e.g., religious beliefs, Y is the outcome of interest, hiring decision, and Z and W are covariates encoding, respectively, the educational level and proximity to work of the applicant.

Causal Model

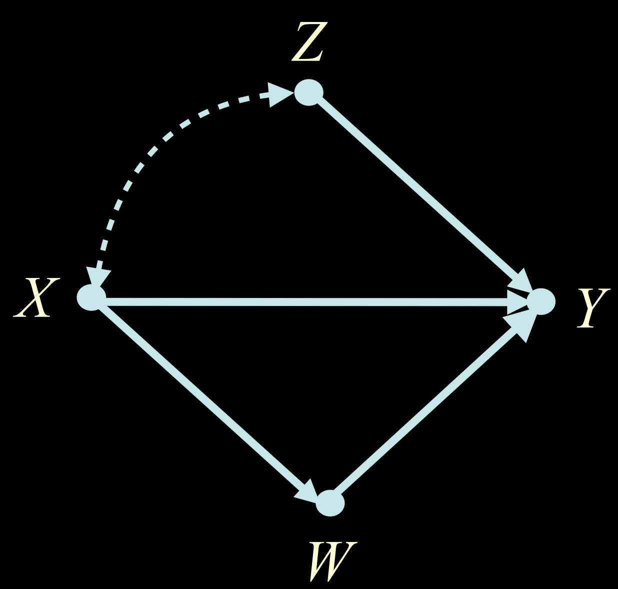

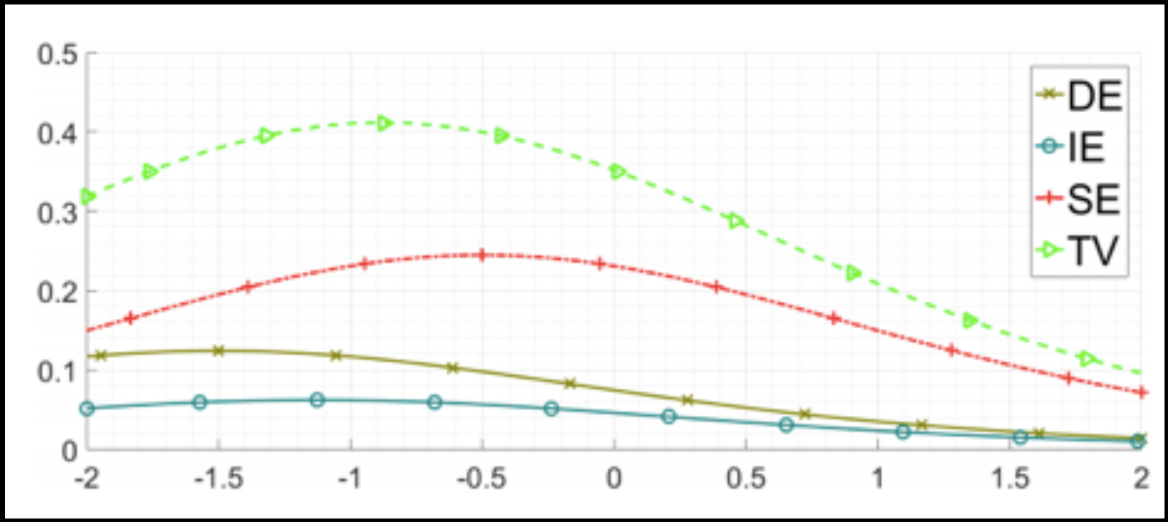

CAUSAL FAIRNESS ANALYSIS

- Example: Assume X is the protected attribute, e.g., religious beliefs, Y is the outcome of interest, hiring decision, and Z and W are covariates encoding, respectively, the educational level and proximity to work of the applicant.

Causal Model

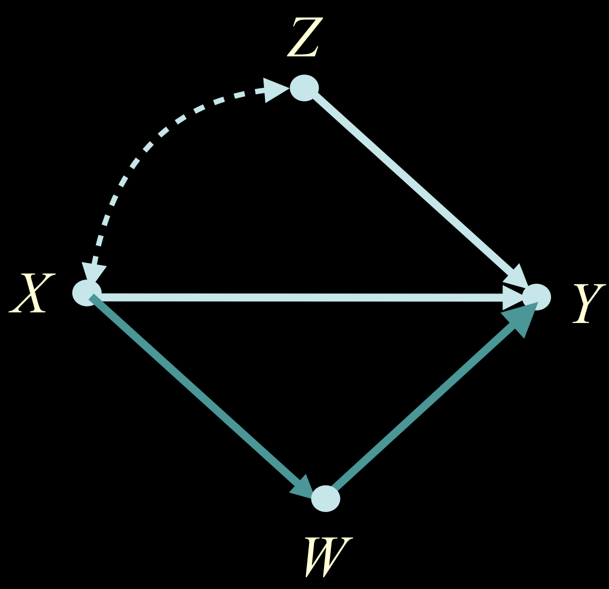

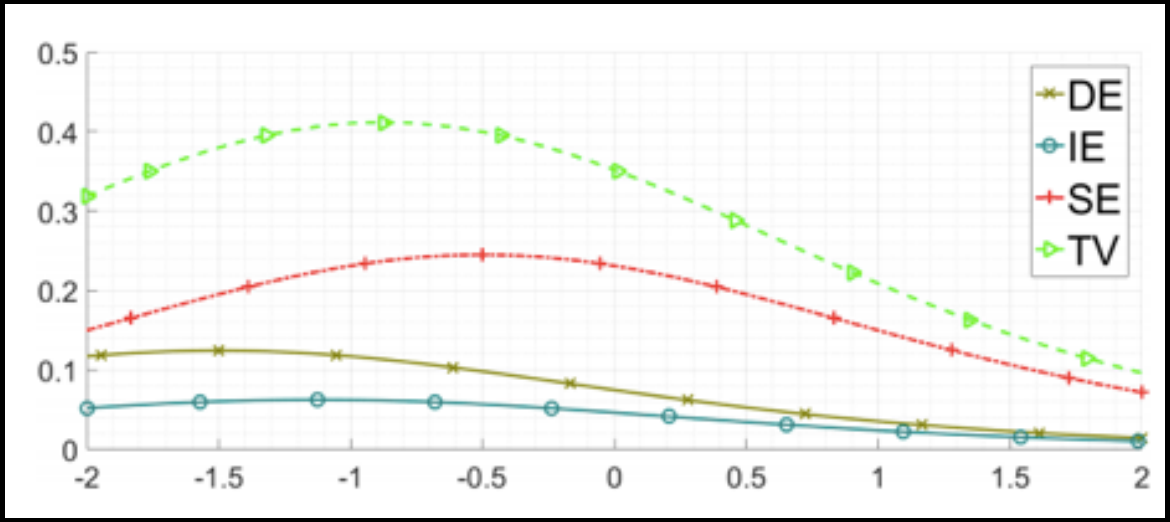

CAUSAL FAIRNESS ANALYSIS

- Example: Assume X is the protected attribute, e.g., religious beliefs, Y is the outcome of interest, hiring decision, and Z and W are covariates encoding, respectively, the educational level and proximity to work of the applicant.

Causal Model

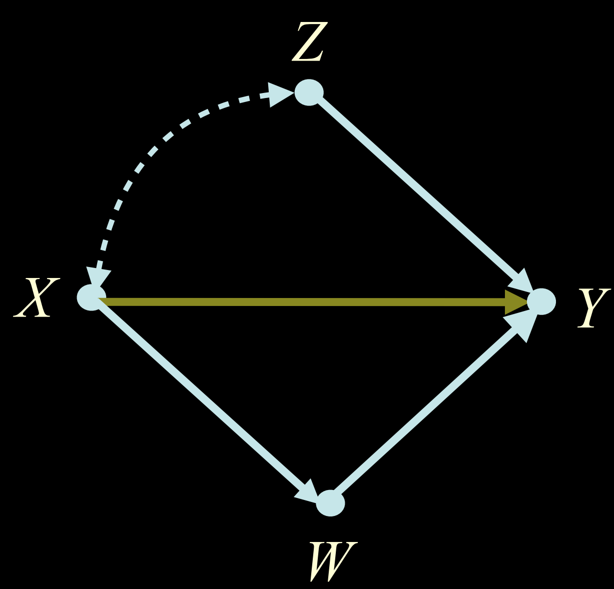

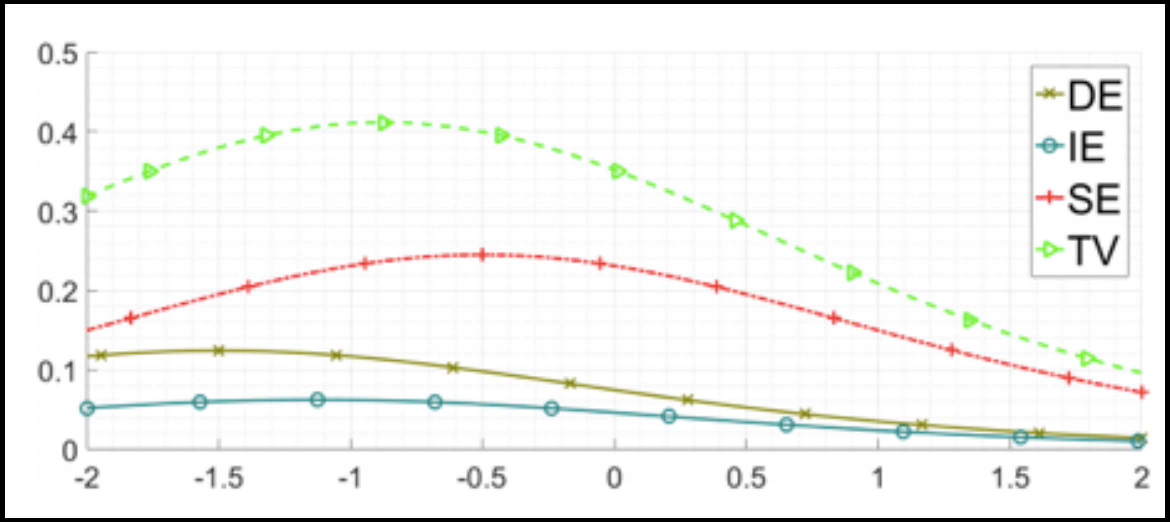

CAUSAL FAIRNESS ANALYSIS

- Example: Assume X is the protected attribute, e.g., religious beliefs, Y is the outcome of interest, hiring decision, and Z and W are covariates encoding, respectively, the educational level and proximity to work of the applicant.

Causal Model