6.7930/HST.956 — Machine Learning for Healthcare
Spring 2025

Recitation 6

Causal Inference

Ignorability Assumption
Sensitivity Analysis / Negative Controls
Overlap Assumption
Extrapolation

Ilker Demirel, 3/7/2025

Potential Outcomes / Identification

What would have happened under treatment T = 0 / T=1?

$$Y(0),\ Y(1)$$

Potential Outcomes / Identification

What would have happened under treatment T = 0 / T=1?

$$Y(0),\ Y(1)$$

Only one PO is observed for each patient:

$$Y = Y(T)$$

"Fundamental problem of causal inference"

X	T	Y	$Y(0)$	$Y(1)$
0 (F)	0	1	1	1
0 (F)	0	1	1	2
0 (F)	0	2	2	1
0 (F)	1	2	1	2
1 (M)	0	2	2	10
1 (M)	1	12	2	12
1 (M)	1	11	1	11
1 (M)	1	10	1	10

Potential Outcomes / Identification

What would have happened under treatment T = 0 / T=1?

Only one PO is observed for each patient: $Y(0),\ Y(1)$

$$Y = Y(T)$$

"Fundamental problem of causal inference"

Average Treatment Effect (ATE):

$$\mathbb{E}[Y(1) - Y(0)]$$

$$= \mathbb{E}[Y(1)] - \mathbb{E}[Y(0)]$$

X	T	Y	$Y(0)$	$Y(1)$
0 (F)	0	1	1	1
0 (F)	0	1	1	2
0 (F)	0	2	2	1
0 (F)	1	2	1	2
1 (M)	0	2	2	10
1 (M)	1	12	2	12
1 (M)	1	11	1	11
1 (M)	1	10	1	10

Potential Outcomes / Identification

What would have happened under treatment T = 0 / T=1?

Only one PO is observed for each patient: $Y(0),\ Y(1)$

$$Y = Y(T)$$

"Fundamental problem of causal inference"

Average Treatment Effect (ATE):

$$\mathbb{E}[Y(1) - Y(0)]$$
$$= \mathbb{E}[Y(1)] - \mathbb{E}[Y(0)]$$
$$= \mathbb{E}[Y(1) \mid T = 1] - \mathbb{E}[Y(0) \mid T = 0]$$

?

X	T	Y	$Y(0)$	$Y(1)$
0 (F)	0	1	1	1
0 (F)	0	1	1	2
0 (F)	0	2	2	1
0 (F)	1	2	1	2
1 (M)	0	2	2	10
1 (M)	1	12	2	12
1 (M)	1	11	1	11
1 (M)	1	10	1	10

Potential Outcomes / Identification

What would have happened under treatment T = 0 / T=1?

$$Y(0),\ Y(1)$$

Only one PO is observed for each patient:

$$Y = Y(T)$$

"Fundamental problem of causal inference"

X	T	Y	$Y(0)$	$Y(1)$
0 (F)	0	1	1	1
0 (F)	0	1	1	2
0 (F)	0	2	2	1
0 (F)	1	2	1	2
1 (M)	0	2	2	10
1 (M)	1	12	2	12
1 (M)	1	11	1	11
1 (M)	1	10	1	10

Average Treatment Effect (ATE):

$$\begin{aligned} & \mathbb{E}[Y(1) - Y(0)] \\ &= \mathbb{E}[Y(1)] - \mathbb{E}[Y(0)] \\ &= \mathbb{E}[Y(1) \mid T = 1] - \mathbb{E}[Y(0) \mid T = 0] \end{aligned}$$

?

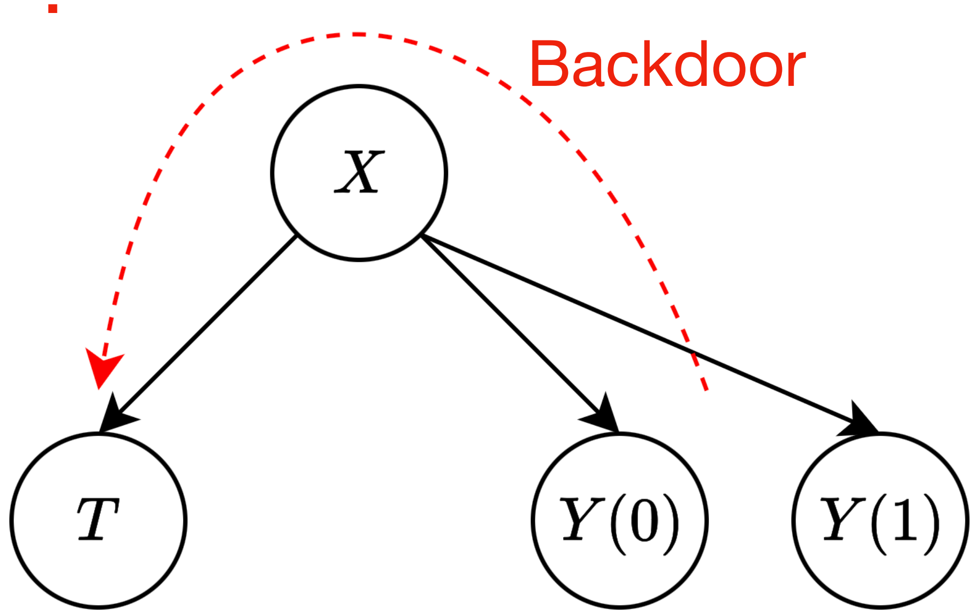

Potential Outcomes / Identification

What would have happened under treatment T = 0 / T=1?

$$Y(0),\ Y(1)$$

Only one PO is observed for each patient:

$$Y = Y(T)$$

"Fundamental problem of causal inference"

X	T	Y	$Y(0)$	$Y(1)$
0 (F)	0	1	1	1
0 (F)	0	1	1	2
0 (F)	0	2	2	1
0 (F)	1	2	1	2
1 (M)	0	2	2	10
1 (M)	1	12	2	12
1 (M)	1	11	1	11
1 (M)	1	10	1	10

Average Treatment Effect (ATE):

$$\mathbb{E}[Y(1) - Y(0)]$$
$$= \mathbb{E}[Y \mid T=1] - \mathbb{E}[Y \mid T=0]$$

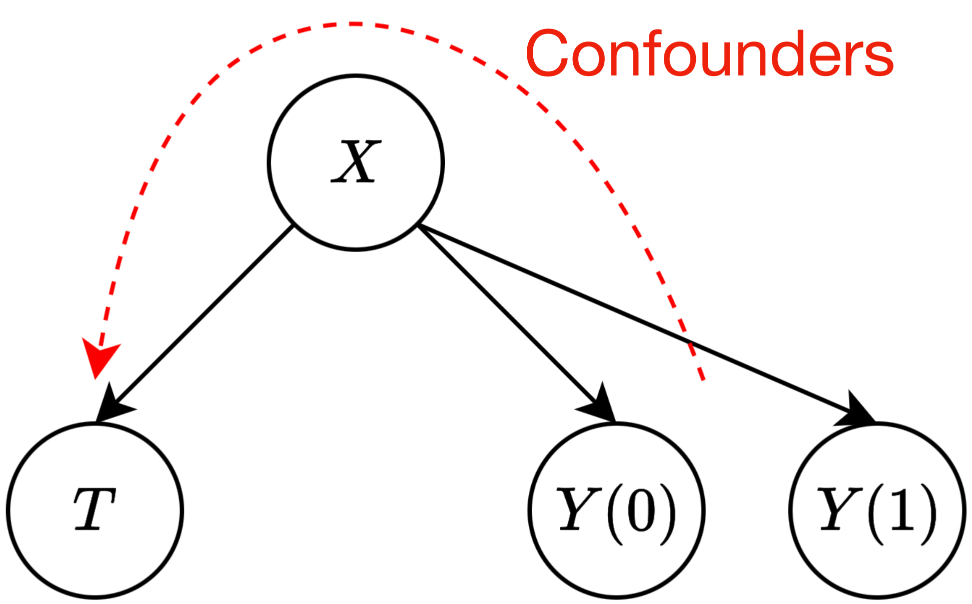

$$\mathbb{E}[Y(1) \mid X] - \mathbb{E}[Y(0) \mid X]$$
$$= \mathbb{E}[Y(1) \mid X, T=1] - \mathbb{E}[Y(0) \mid X, T=0]$$
$$= \mathbb{E}_X\big[\mathbb{E}[Y \mid X, T=1] - \mathbb{E}[Y \mid X, T=0]\big]$$

Potential Outcomes / Identification

What would have happened under treatment T = 0 / T=1?

Only one PO is observed for each patient: $Y(0),\ Y(1)$

$$Y = Y(T)$$

"Fundamental problem of causal inference"

X	T	Y	$Y(0)$	$Y(1)$
0 (F)	0	1	1	1
0 (F)	0	1	1	2
0 (F)	0	2	2	1
0 (F)	1	2	1	2
1 (M)	0	2	2	10
1 (M)	1	12	2	12
1 (M)	1	11	1	11
1 (M)	1	10	1	10

Average Treatment Effect (ATE):

$$\mathbb{E}[Y(1) - Y(0)]$$
$$= \mathbb{E}[Y(1)] - \mathbb{E}[Y(0)]$$
$$= \mathbb{E}[Y(1) \,|\, T = 1] - \mathbb{E}[Y(0) \,|\, T = 0]$$

?

$$Y(0), Y(1) \perp\!\!\!\perp T \mid X$$

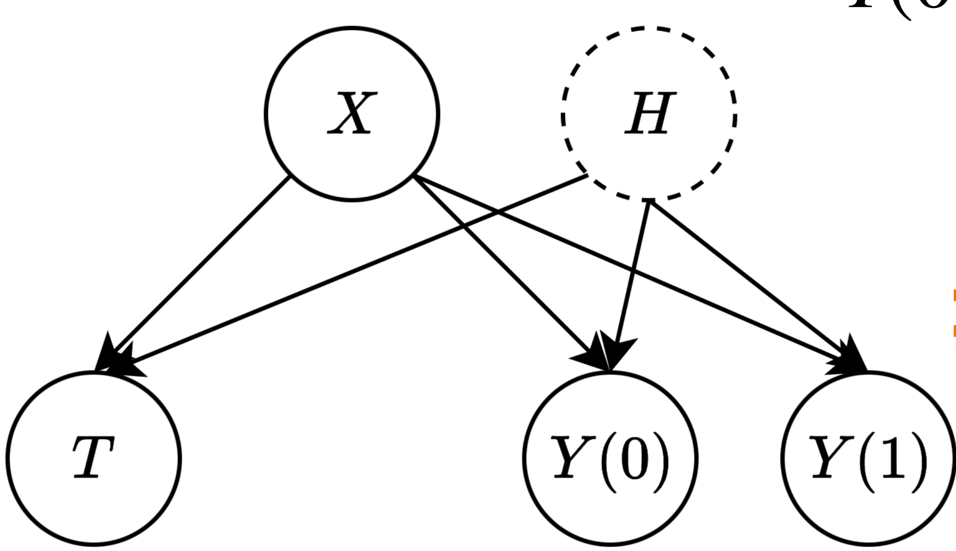

Ignorability of treatment assignment
ASSUMPTION

$$\Longrightarrow$$

$$\mathbb{E}[Y(1) \mid X] - \mathbb{E}[Y(0) \mid X]$$
$$= \mathbb{E}[Y(1) \,|\, X, T = 1] - \mathbb{E}[Y(0) \,|\, X, T = 0]$$
$$= \mathbb{E}[Y \,|\, X, T = 1] - \mathbb{E}[Y \,|\, X, T = 0]$$

Identification

Checking Ignorability Assumption

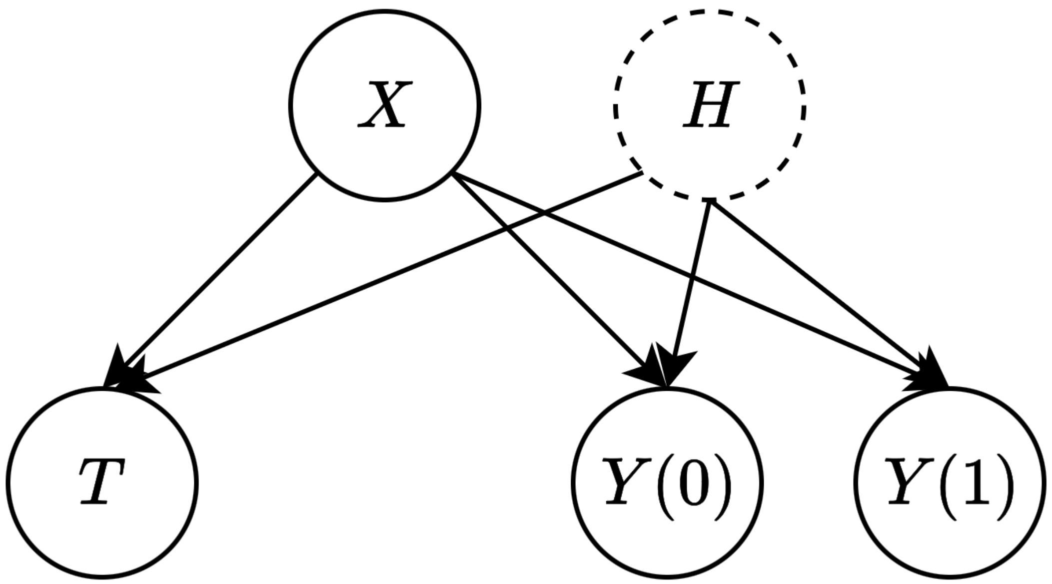

X can be very high dim.

- socio-economic status (SES)
- comorbidities
- past complications
- genetics
- demographics
- imaging data
- etc…

Checking Ignorability Assumption

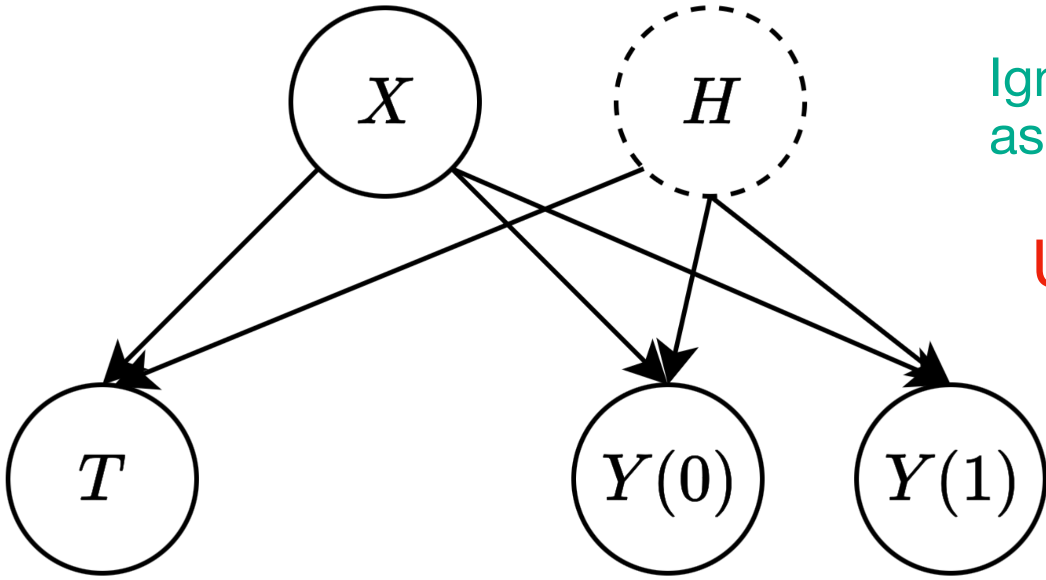

Ignorability of treatment assignment assumption

$$Y(0), Y(1) \perp\!\!\!\perp T \mid X$$

Unverifiable, BUT there are things you can do

X can be very high dim.

- socio-economic status (SES)
- comorbidities
- past complications
- genetics
- demographics
- imaging data
- etc…

Checking Ignorability Assumption

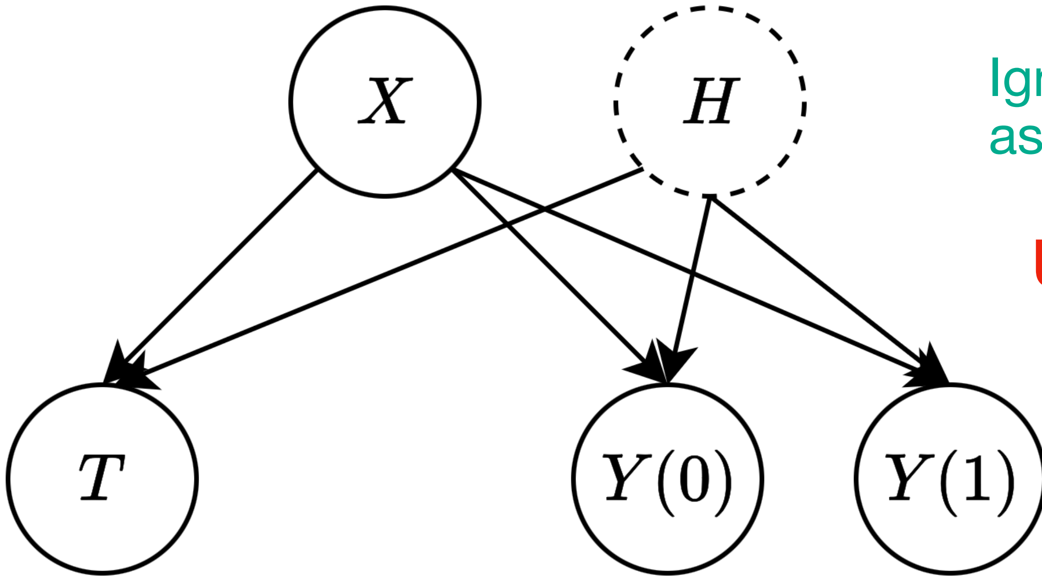

Ignorability of treatment assignment assumption

$$Y(0), Y(1) \perp\!\!\!\perp T \mid X$$

Unverifiable, BUT there are things you can do

Sensitivity Analysis: How strong would an unmeasured/hidden factor need to be to explain away our findings (treatment effect estimate) ?

X can be very high dim.

- socio-economic status (SES)
- comorbidities
- past complications
- genetics
- demographics
- imaging data
- etc…

Checking Ignorability Assumption

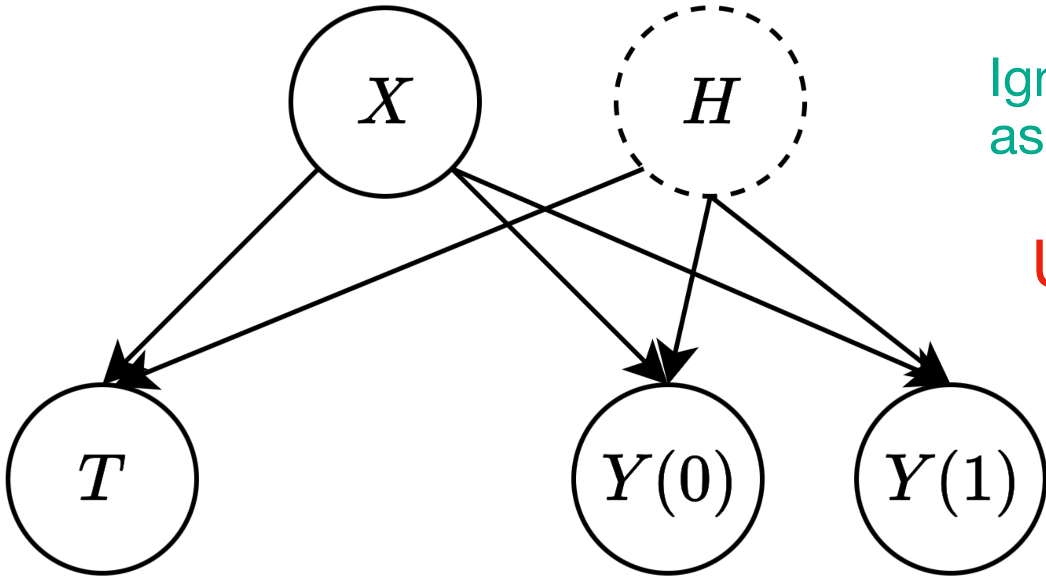

Ignorability of treatment assignment assumption

$$Y(0), Y(1) \perp\!\!\!\perp T \mid X$$

Unverifiable, BUT there are things you can do

X can be very high dim.

- socio-economic status (SES)
- comorbidities
- past complications
- genetics
- demographics
- imaging data
- etc…

Sensitivity Analysis: How strong would an unmeasured/hidden factor need to be to explain away our findings (treatment effect estimate) ?

Example Output: For the treatment effect to be negative, H would need to make receiving the treatment 1.5x likely AND experiencing a positive outcome 2x likely.

Checking Ignorability Assumption

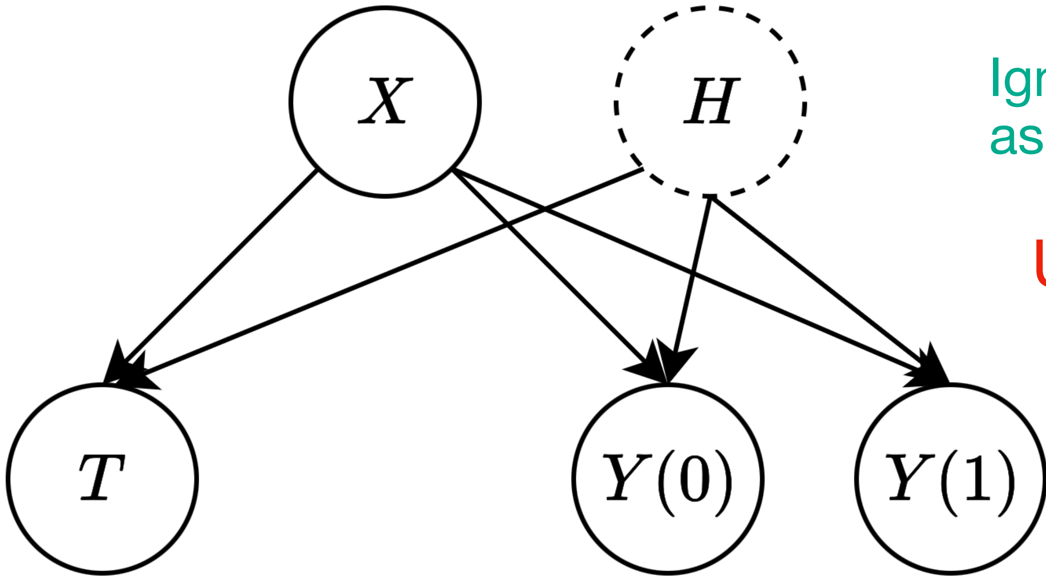

X can be very high dim.

- socio-economic status (SES)
- comorbidities
- past complications
- genetics
- demographics
- imaging data
- etc…

Ignorability of treatment assignment assumption

$$Y(0), Y(1) \perp\!\!\!\perp T \mid X$$

Unverifiable, BUT there are things you can do

Sensitivity Analysis: How strong would an unmeasured/hidden factor need to be to explain away our findings (treatment effect estimate) ?

Example Output: For the treatment effect to be negative, H would need to make receiving the treatment 1.5x likely AND experiencing a positive outcome 2x likely.

Drawback: Not very interpretable/actionable. Depends on other assumptions itself.

Checking Ignorability Assumption

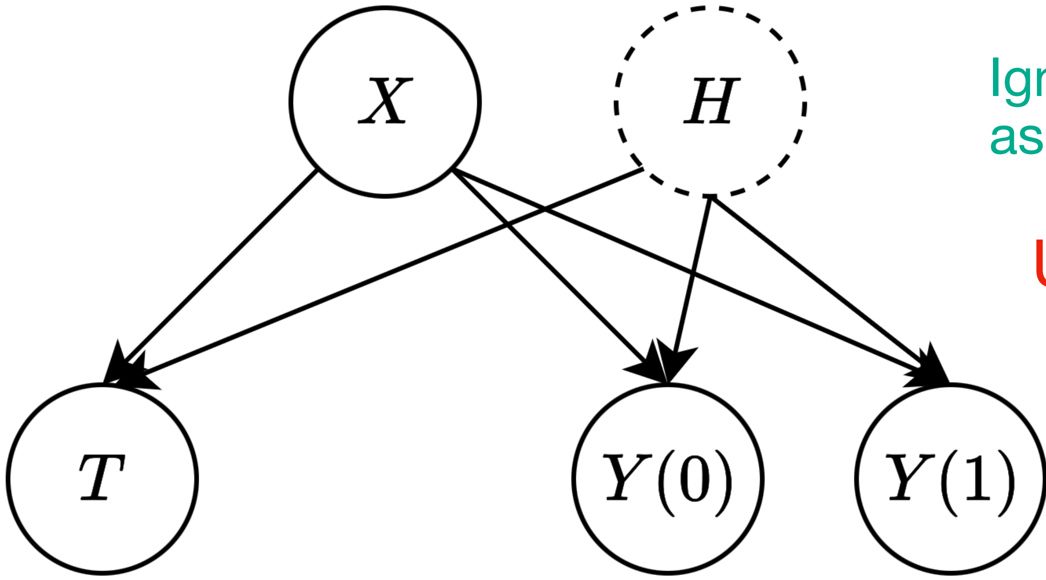

Ignorability of treatment assignment assumption

$$Y(0), Y(1) \perp\!\!\!\perp T \mid X$$

Unverifiable, BUT there are things you can do

X can be very high dim.

- socio-economic status (SES)
- comorbidities
- past complications
- genetics
- demographics
- imaging data
- etc…

Sensitivity Analysis: How strong would an unmeasured/hidden factor need to be to explain away our findings (treatment effect estimate) ?

Example Output: For the treatment effect to be negative, H would need to make receiving the treatment 1.5x likely AND experiencing a positive outcome 2x likely.

Drawback: Not very interpretable/actionable. Depends on other assumptions itself.

Negative controls: Something we can use for a sanity check.

E.g.: An outcome we know the treatment should not have any effect on.

Negative Controls

The NEW ENGLAND JOURNAL of MEDICINE

https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

Negative Controls

The NEW ENGLAND JOURNAL of MEDICINE
https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

Table 1. (Continued.)

Characteristics	Unvaccinated Controls (N=596,618)	Vaccinated Persons (N=596,618)
Cerebrovascular disease	18,392 (3.1)	17,792 (3.0)
Other respiratory disease	2,198 (0.4)	2,014 (0.3)
Hypertension	101,017 (16.9)	103,028 (17.3)
Immunosuppression	15,823 (2.7)	16,180 (2.7)
Neurologic disease	25,897 (4.3)	24,111 (4.0)
Liver disease	11,109 (1.9)	9,699 (1.6)
Overweight: BMI, 25 to 30	203,296 (34.1)	212,778 (35.7)
Thalassemia	3,764 (0.6)	3,967 (0.7)
Type 1 diabetes mellitus	2,309 (0.4)	2,406 (0.4)

Negative Controls

The NEW ENGLAND JOURNAL of MEDICINE

https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

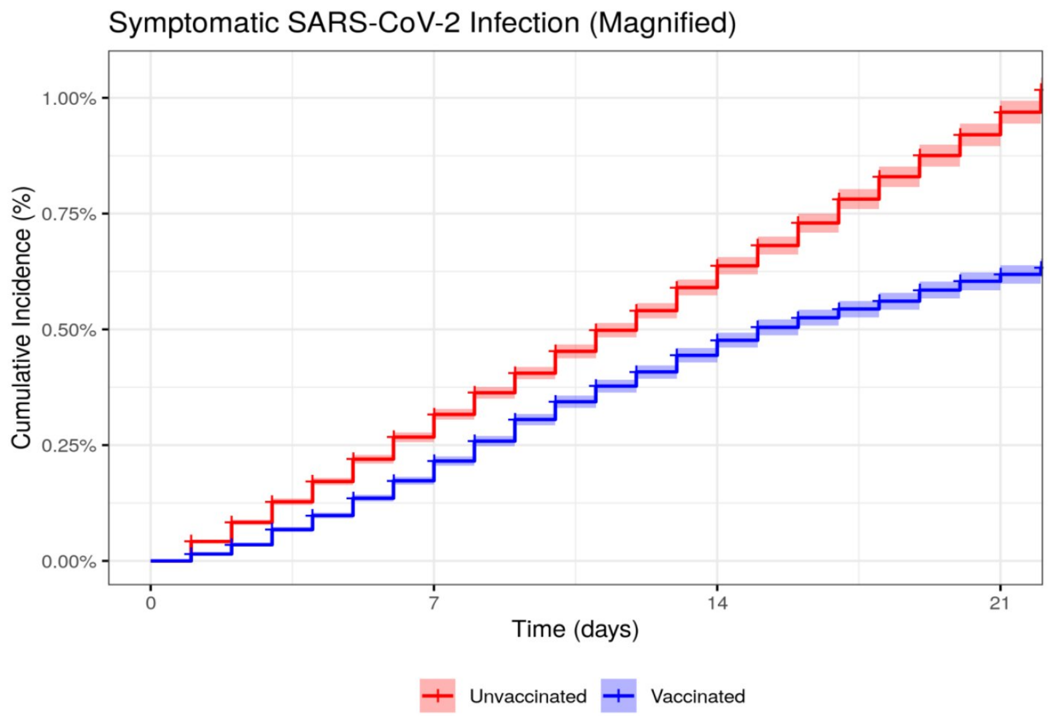

Table 1. (Continued.)

Characteristics	Unvaccinated Controls (N=596,618)	Vaccinated Persons (N=596,618)
Cerebrovascular disease	18,392 (3.1)	17,792 (3.0)
Other respiratory disease	2,198 (0.4)	2,014 (0.3)
Hypertension	101,017 (16.9)	103,028 (17.3)
Immunosuppression	15,823 (2.7)	16,180 (2.7)
Neurologic disease	25,897 (4.3)	24,111 (4.0)
Liver disease	11,109 (1.9)	9,699 (1.6)
Overweight: BMI, 25 to 30	203,296 (34.1)	212,778 (35.7)
Thalassemia	3,764 (0.6)	3,967 (0.7)
Type 1 diabetes mellitus	2,309 (0.4)	2,406 (0.4)

Negative Controls

The NEW ENGLAND JOURNAL of MEDICINE

https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

Symptomatic SARS-CoV-2 Infection (Magnified)

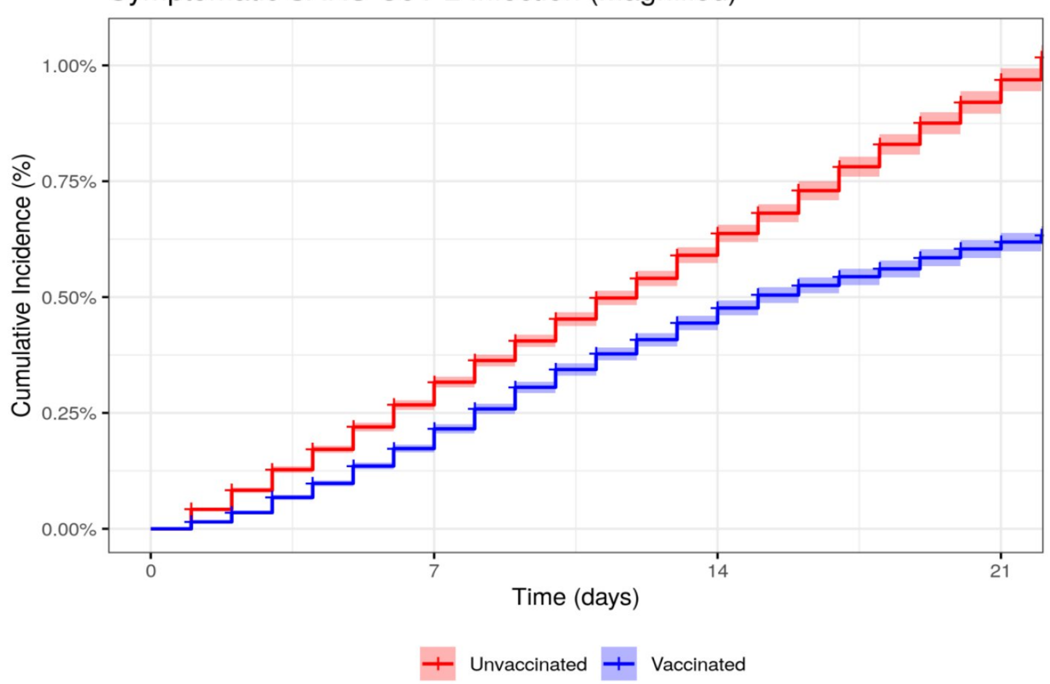

Miguel Hernán @_MiguelHernan · Feb 24, 2021

5/
No, it doesn't.

After matching on age (and sex), the curves of infection start to diverge from day 0, which indicates that the vaccinated had a lower risk of infection than the unvaccinated.

https://x.com/_MiguelHernan/status/1364700315044438023

Table 1. **(Continued.)**

Characteristics	Unvaccinated Controls (N=596,618)	Vaccinated Persons (N=596,618)
Cerebrovascular disease	18,392 (3.1)	17,792 (3.0)
Other respiratory disease	2,198 (0.4)	2,014 (0.3)
Hypertension	101,017 (16.9)	103,028 (17.3)
Immunosuppression	15,823 (2.7)	16,180 (2.7)
Neurologic disease	25,897 (4.3)	24,111 (4.0)
Liver disease	11,109 (1.9)	9,699 (1.6)
Overweight: BMI, 25 to 30	203,296 (34.1)	212,778 (35.7)
Thalassemia	3,764 (0.6)	3,967 (0.7)
Type 1 diabetes mellitus	2,309 (0.4)	2,406 (0.4)

Negative Controls

The NEW ENGLAND JOURNAL of MEDICINE

https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

Symptomatic SARS-CoV-2 Infection (Magnified)

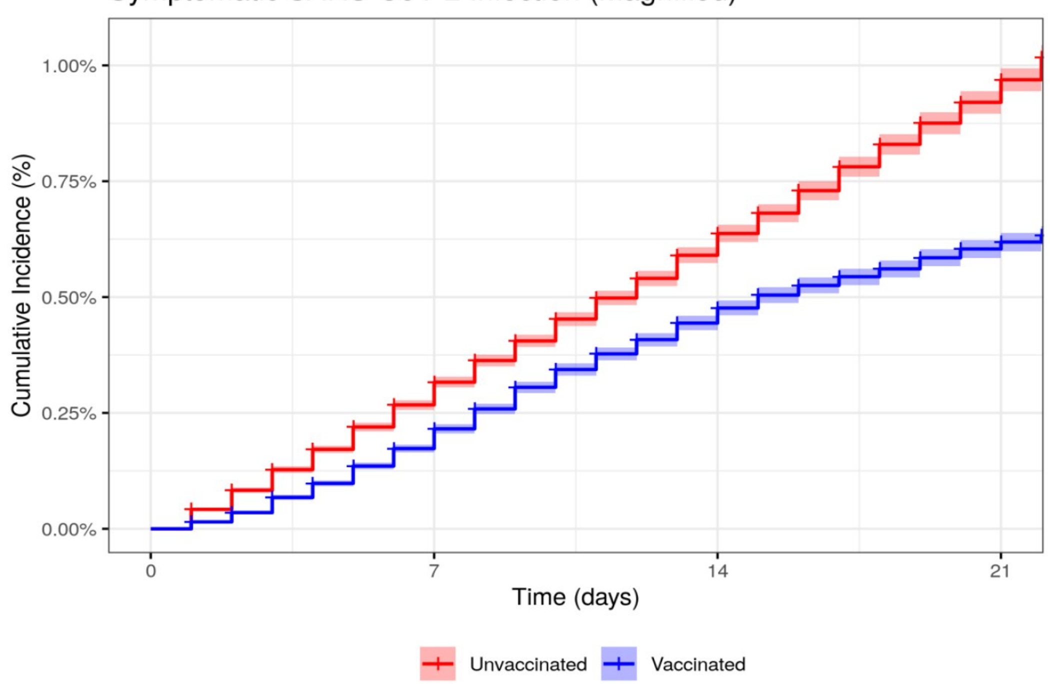

Table 1. (Continued.)

Characteristics	Unvaccinated Controls (N=596,618)	Vaccinated Persons (N=596,618)
Cerebrovascular disease	18,392 (3.1)	17,792 (3.0)
Other respiratory disease	2,198 (0.4)	2,014 (0.3)
Hypertension	101,017 (16.9)	103,028 (17.3)
Immunosuppression	15,823 (2.7)	16,180 (2.7)
Neurologic disease	25,897 (4.3)	24,111 (4.0)
Liver disease	11,109 (1.9)	9,699 (1.6)
Overweight: BMI, 25 to 30	203,296 (34.1)	212,778 (35.7)
Thalassemia	3,764 (0.6)	3,967 (0.7)
Type 1 diabetes mellitus	2,309 (0.4)	2,406 (0.4)

A Documented SARS-CoV-2 Infection

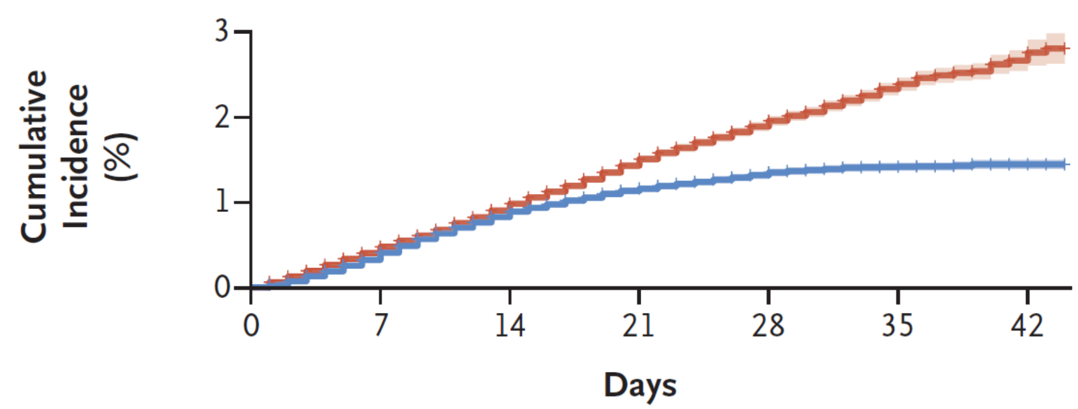

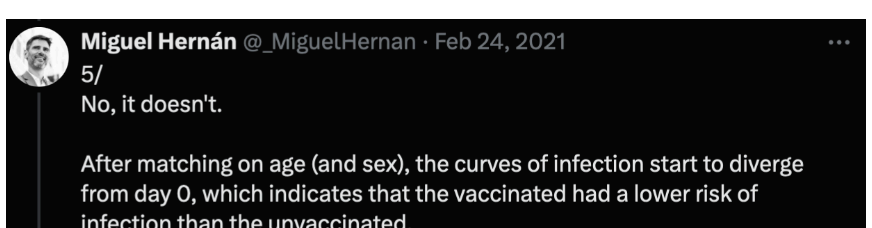
Miguel Hernán @_MiguelHernan · Feb 24, 2021
5/
No, it doesn't.

After matching on age (and sex), the curves of infection start to diverge from day 0, which indicates that the vaccinated had a lower risk of infection than the unvaccinated.

https://x.com/_MiguelHernan/status/1364700315044438023

Negative Controls

The NEW ENGLAND JOURNAL *of* MEDICINE

https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

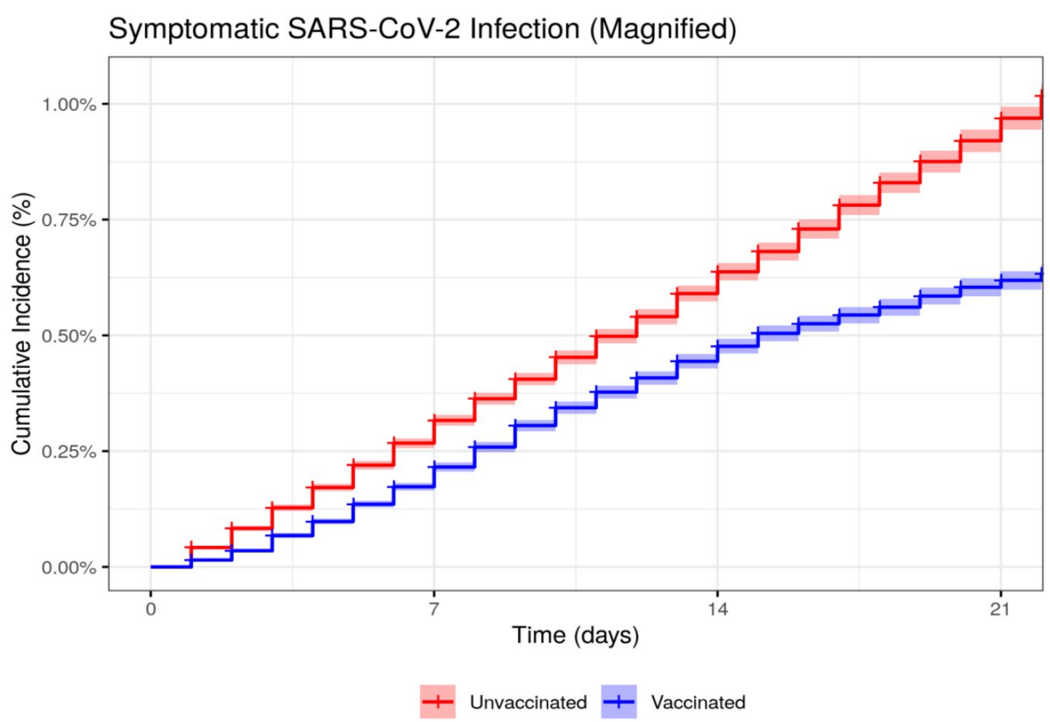

Table 1. (Continued.)

Characteristics	Unvaccinated Controls (N=596,618)	Vaccinated Persons (N=596,618)
Cerebrovascular disease	18,392 (3.1)	17,792 (3.0)
Other respiratory disease	2,198 (0.4)	2,014 (0.3)
Hypertension	101,017 (16.9)	103,028 (17.3)
Immunosuppression	15,823 (2.7)	16,180 (2.7)
Neurologic disease	25,897 (4.3)	24,111 (4.0)
Liver disease	11,109 (1.9)	9,699 (1.6)
Overweight: BMI, 25 to 30	203,296 (34.1)	212,778 (35.7)
Thalassemia	3,764 (0.6)	3,967 (0.7)
Type 1 diabetes mellitus	2,309 (0.4)	2,406 (0.4)

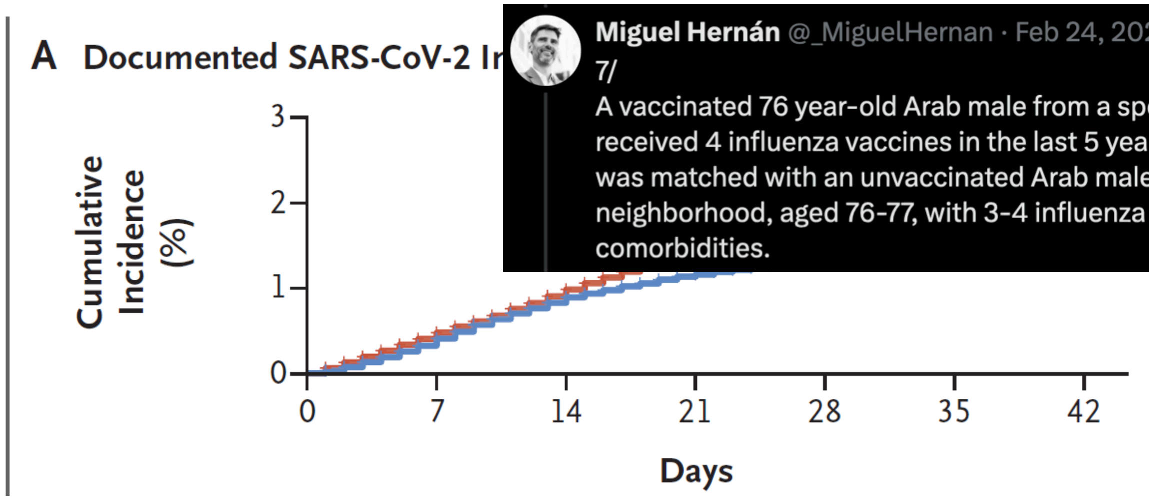

Miguel Hernán @_MiguelHernan · Feb 24, 2021

7/

A vaccinated 76 year-old Arab male from a specific neighborhood who received 4 influenza vaccines in the last 5 years and had 2 comorbidities was matched with an unvaccinated Arab male from the same neighborhood, aged 76-77, with 3-4 influenza vaccines and 2 comorbidities.

Miguel Hernán @_MiguelHernan · Feb 24, 2021

5/

No, it doesn't.

After matching on age (and sex), the curves of infection start to diverge from day 0, which indicates that the vaccinated had a lower risk of infection than the unvaccinated.

https://x.com/_MiguelHernan/status/1364700315044438023

Negative Controls

The NEW ENGLAND JOURNAL *of* MEDICINE

https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

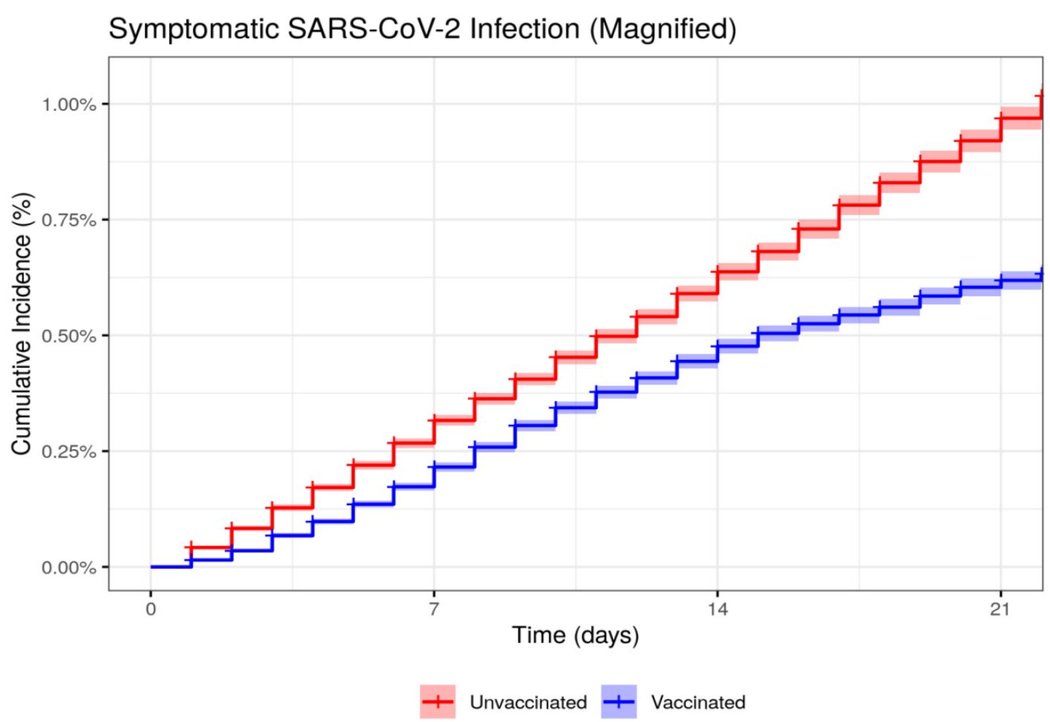

Table 1. (Continued.)

Characteristics	Unvaccinated Controls (N=596,618)	Vaccinated Persons (N=596,618)
Cerebrovascular disease	18,392 (3.1)	17,792 (3.0)
Other respiratory disease	2,198 (0.4)	2,014 (0.3)
Hypertension	101,017 (16.9)	103,028 (17.3)
Immunosuppression	15,823 (2.7)	16,180 (2.7)
Neurologic disease	25,897 (4.3)	24,111 (4.0)
Liver disease	11,109 (1.9)	9,699 (1.6)
Overweight: BMI, 25 to 30	203,296 (34.1)	212,778 (35.7)
Thalassemia	3,764 (0.6)	3,967 (0.7)
Type 1 diabetes mellitus	2,309 (0.4)	2,406 (0.4)

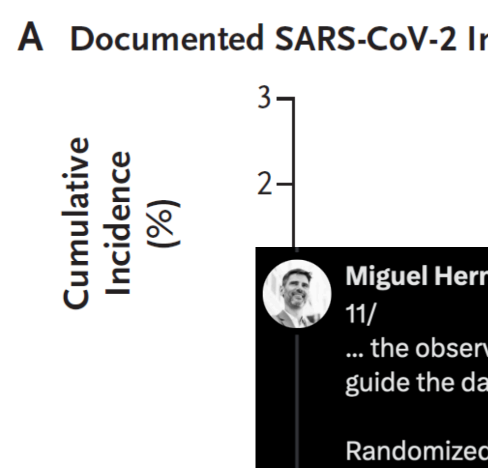

Miguel Hernán @_MiguelHernan · Feb 24, 2021

7/
A vaccinated 76 year-old Arab male from a specific neighborhood who received 4 influenza vaccines in the last 5 years and had 2 comorbidities was matched with an unvaccinated Arab male from the same neighborhood, aged 76-77, with 3-4 influenza vaccines and 2 comorbidities.

Miguel Hernán @_MiguelHernan · Feb 24, 2021

11/
... the observational study needs the trial's findings as a #benchmark to guide the data analysis and strengthen the quality of the #causalinference.

Randomized trials & Observational studies working together. The best of both worlds.

Let's keep doing it after the pandemic

Miguel Hernán @_MiguelHernan · Feb 24, 2021

5/
No, it doesn't.

After matching on age (and sex), the curves of infection start to diverge from day 0, which indicates that the vaccinated had a lower risk of infection than the unvaccinated.

https://x.com/_MiguelHernan/status/1364700315044438023

Negative Controls

The NEW ENGLAND JOURNAL *of* MEDICINE

https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

Symptomatic SARS-CoV-2 Infection (Magnified)

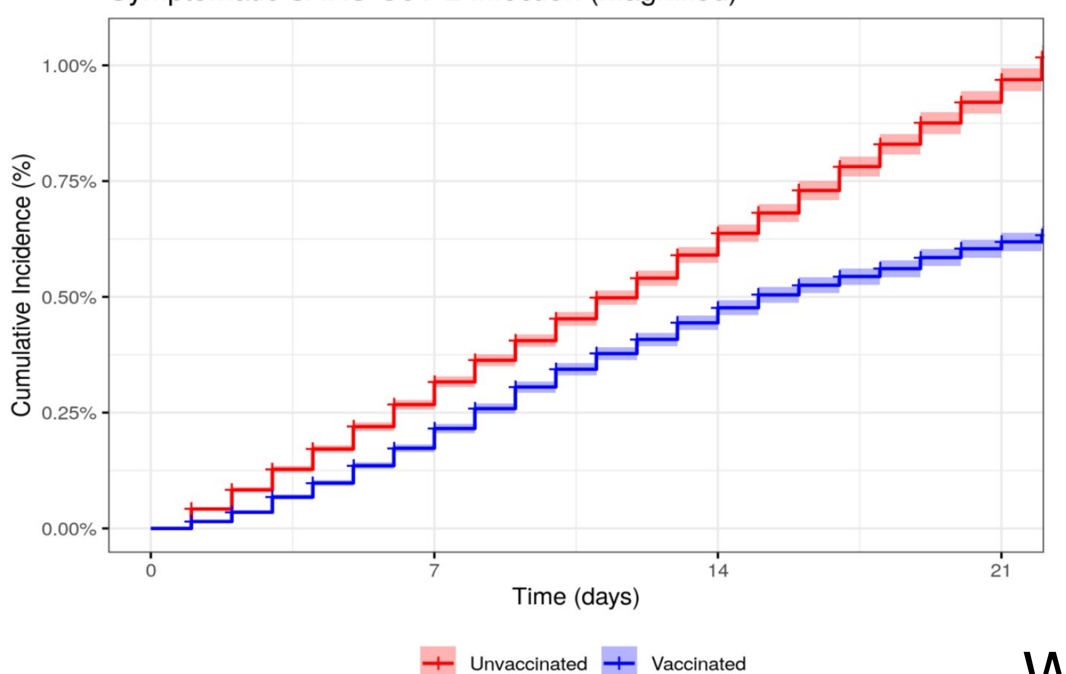

Table 1. (Continued.)

Characteristics	Unvaccinated Controls (N=596,618)	Vaccinated Persons (N=596,618)
Cerebrovascular disease	18,392 (3.1)	17,792 (3.0)
Other respiratory disease	2,198 (0.4)	2,014 (0.3)
Hypertension	101,017 (16.9)	103,028 (17.3)
Immunosuppression	15,823 (2.7)	16,180 (2.7)
Neurologic disease	25,897 (4.3)	24,111 (4.0)
Liver disease	11,109 (1.9)	9,699 (1.6)
Overweight: BMI, 25 to 30	203,296 (34.1)	212,778 (35.7)
Thalassemia	3,764 (0.6)	3,967 (0.7)
Type 1 diabetes mellitus	2,309 (0.4)	2,406 (0.4)

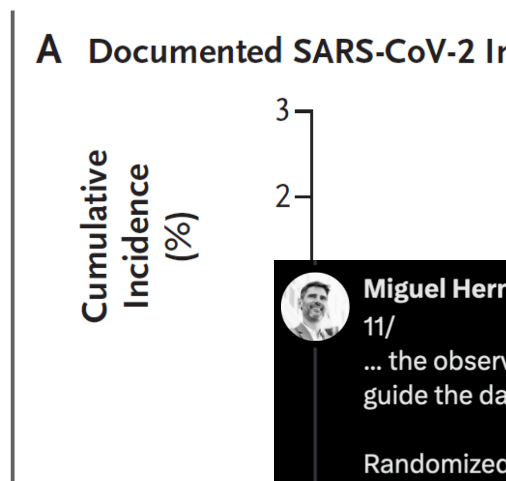

Miguel Hernán @_MiguelHernan · Feb 24, 2021

7/
A vaccinated 76 year-old Arab male from a specific neighborhood who received 4 influenza vaccines in the last 5 years and had 2 comorbidities was matched with an unvaccinated Arab male from the same neighborhood, aged 76-77, with 3-4 influenza vaccines and 2 comorbidities.

Miguel Hernán @_MiguelHernan · Feb 24, 2021

11/
... the observational study needs the trial's findings as a #benchmark to guide the data analysis and strengthen the quality of the #causalinference.

Randomized trials & Observational studies working together. The best of both worlds.

Let's keep doing it after the pandemic

Why not match on more stuff to begin with?

Miguel Hernán @_MiguelHernan · Feb 24, 2021

5/
No, it doesn't.

After matching on age (and sex), the curves of infection start to diverge from day 0, which indicates that the vaccinated had a lower risk of infection than the unvaccinated.

https://x.com/_MiguelHernan/status/1364700315044438023

Negative Controls

The NEW ENGLAND JOURNAL of MEDICINE

https://www.nejm.org/doi/full/10.1056/NEJMoa2101765

ORIGINAL ARTICLE

BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting

Noa Dagan, M.D., Noam Barda, M.D., Eldad Kepten, Ph.D., Oren Miron, M.A., Shay Perchik, M.A., Mark A. Katz, M.D., Miguel A. Hernán, M.D., Marc Lipsitch, D.Phil., Ben Reis, Ph.D., and Ran D. Balicer, M.D.

Symptomatic SARS-CoV-2 Infection (Magnified)

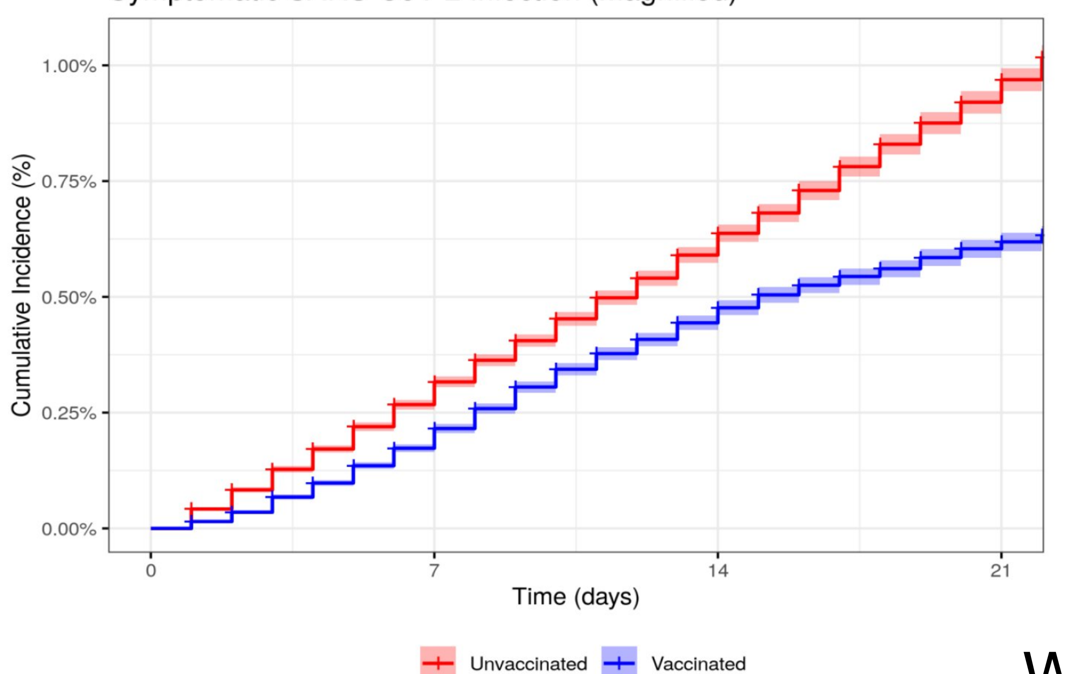

Miguel Hernán @_MiguelHernan · Feb 24, 2021

5/

No, it doesn't.

After matching on age (and sex), the curves of infection start to diverge from day 0, which indicates that the vaccinated had a lower risk of infection than the unvaccinated.

https://x.com/_MiguelHernan/status/1364700315044438023

Table 1. (Continued.)

Characteristics	Unvaccinated Controls (N=596,618)	Vaccinated Persons (N=596,618)
Cerebrovascular disease	18,392 (3.1)	17,792 (3.0)
Other respiratory disease	2,198 (0.4)	2,014 (0.3)
Hypertension	101,017 (16.9)	103,028 (17.3)
Immunosuppression	15,823 (2.7)	16,180 (2.7)
Neurologic disease	25,897 (4.3)	24,111 (4.0)
Liver disease	11,109 (1.9)	9,699 (1.6)
Overweight: BMI, 25 to 30	203,296 (34.1)	212,778 (35.7)
Thalassemia	3,764 (0.6)	3,967 (0.7)
Type 1 diabetes mellitus	2,309 (0.4)	2,406 (0.4)

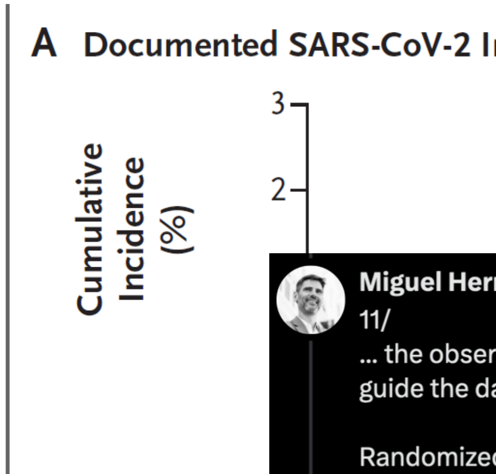

Miguel Hernán @_MiguelHernan · Feb 24, 2021

7/

A vaccinated 76 year-old Arab male from a specific neighborhood who received 4 influenza vaccines in the last 5 years and had 2 comorbidities was matched with an unvaccinated Arab male from the same neighborhood, aged 76-77, with 3-4 influenza vaccines and 2 comorbidities.

Miguel Hernán @_MiguelHernan · Feb 24, 2021

11/

… the observational study needs the trial's findings as a #benchmark to guide the data analysis and strengthen the quality of the #causalinference.

Randomized trials & Observational studies working together. The best of both worlds.

Let's keep doing it after the pandemic

Why not match on more stuff to begin with?
No free lunch. Trade-off:

- Harder to match individuals
- Smaller sample size
- Reduced statistical power in estimates

Overlap

$f_0(X) = X^2 + \epsilon$

$f_1(X) = X^2 + X + \epsilon$

$\text{CATE}(X) = X$

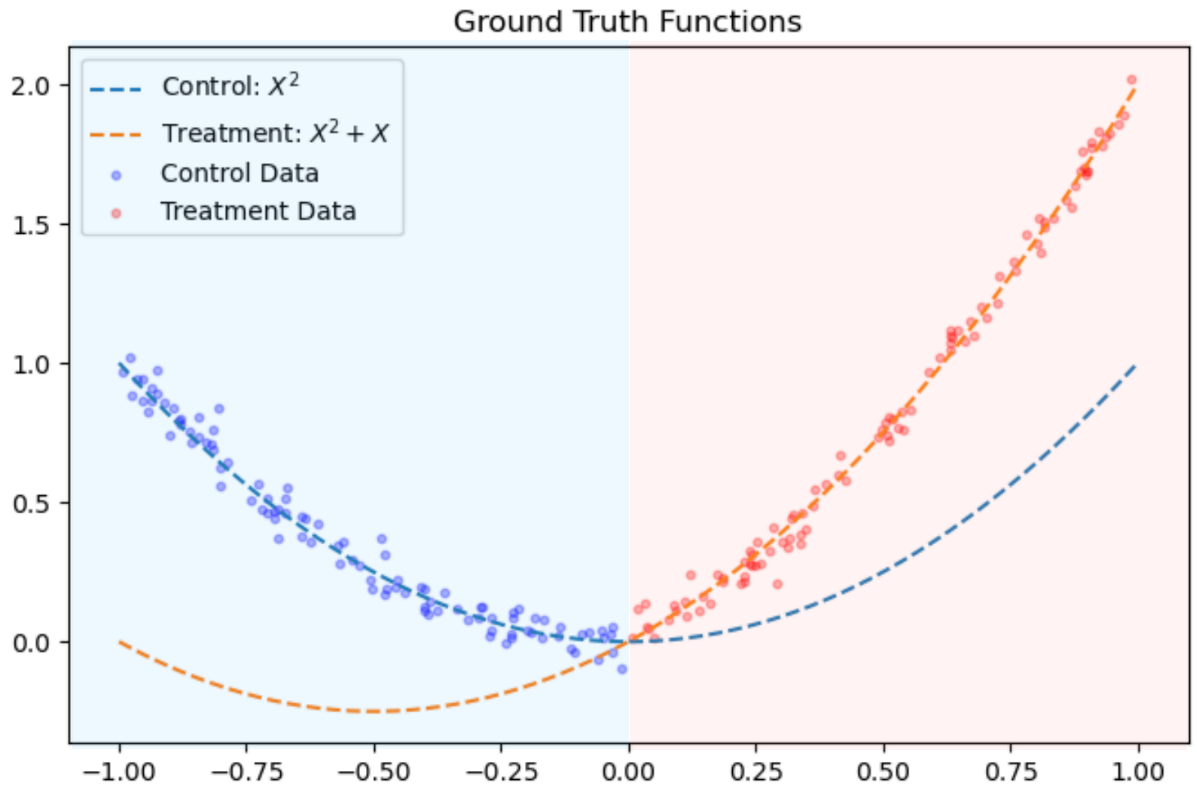

Overlap – Risks of extrapolation

$f_0(X) = X^2 + \epsilon$ $\qquad$ $f_1(X) = X^2 + X + \epsilon$ $\qquad$ $\text{CATE}(X) = X$

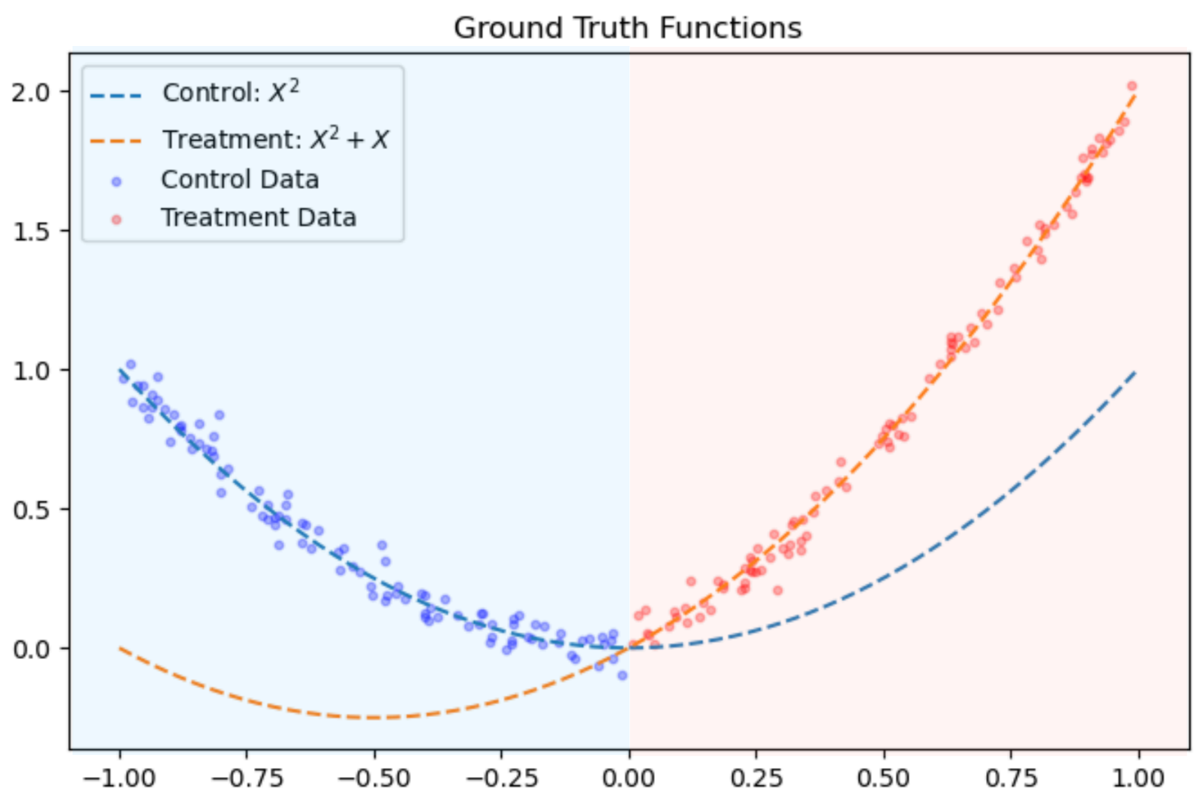

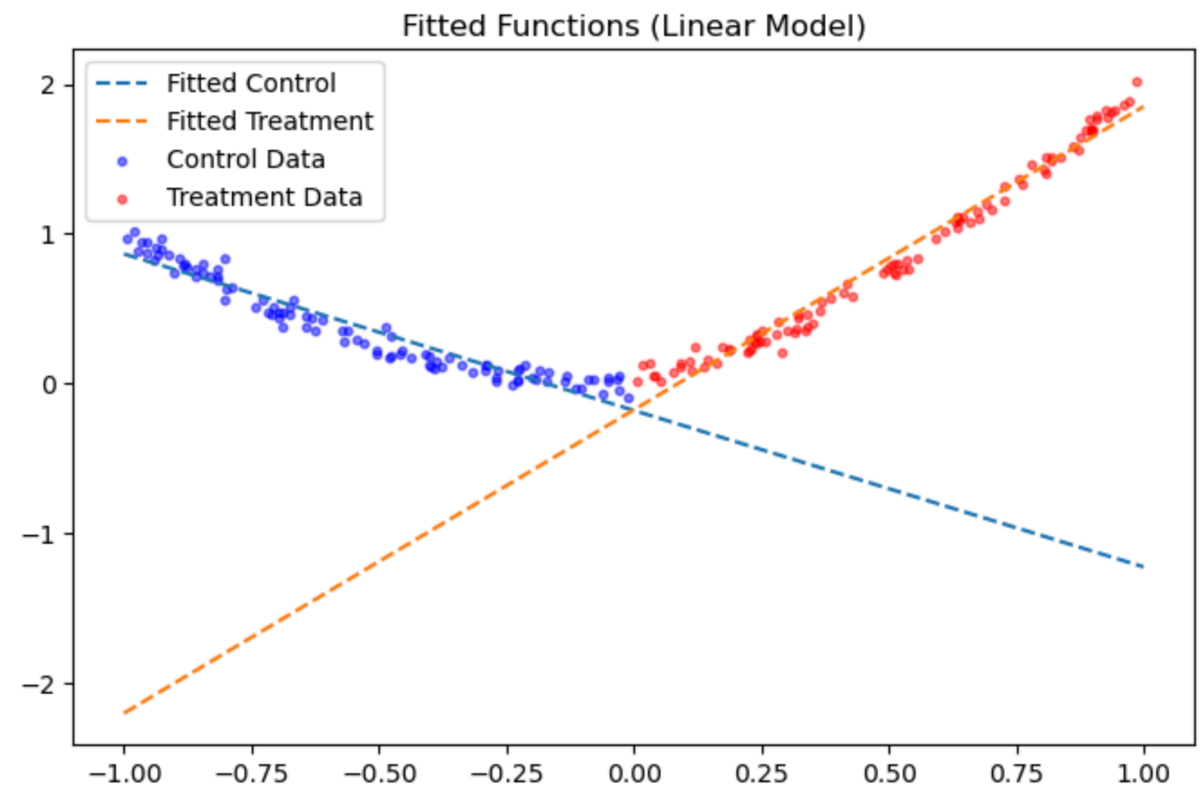

Overlap – Risks of extrapolation

$f_0(X) = X^2 + \epsilon$ $\quad$ $f_1(X) = X^2 + X + \epsilon$ $\quad$ $\text{CATE}(X) = X$

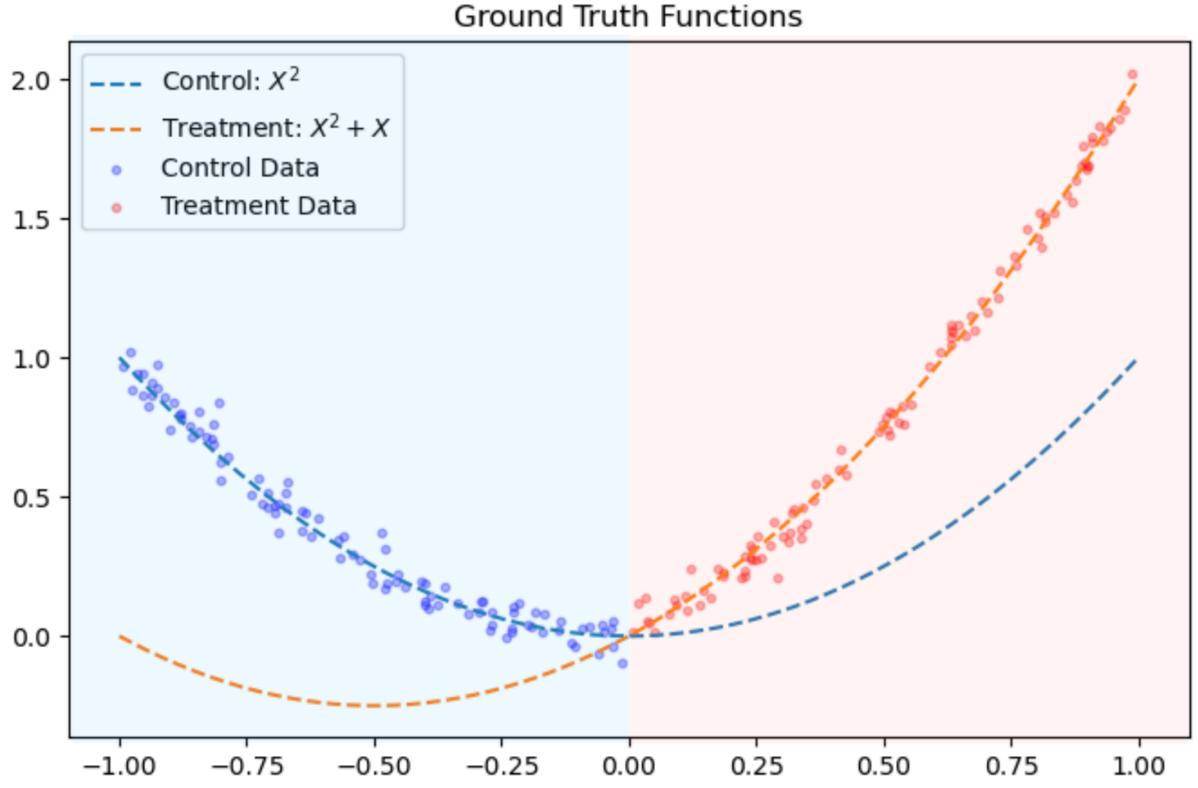

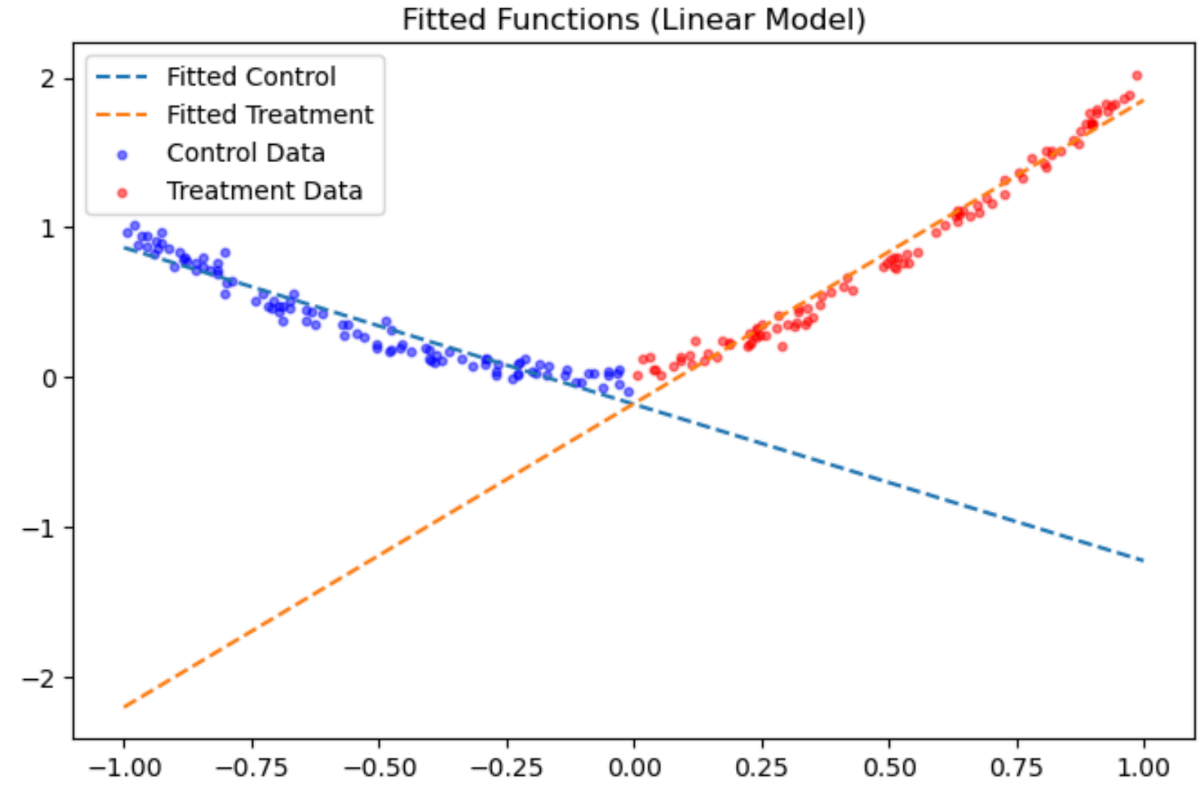

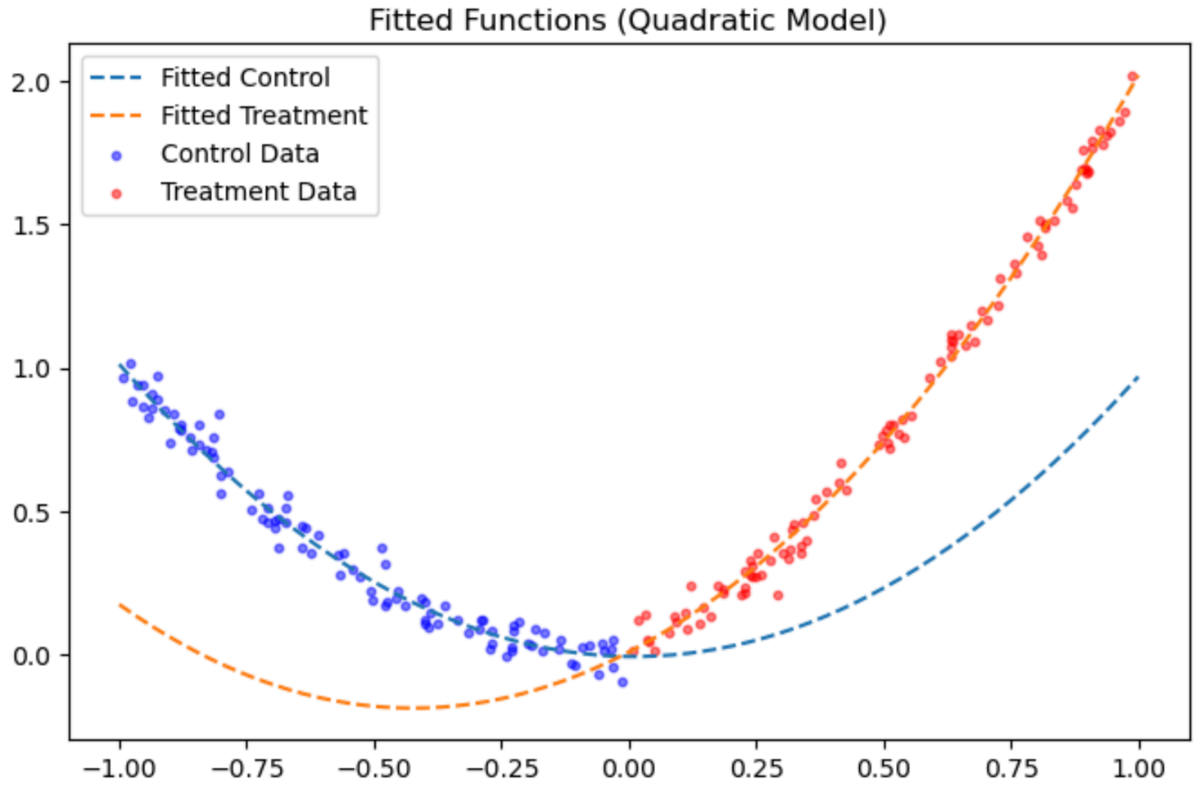

Overlap – Risks of extrapolation

$f_0(X) = X^2 + \epsilon$ $\qquad$ $f_1(X) = X^2 + X + \epsilon$ $\qquad$ $\text{CATE}(X) = X$

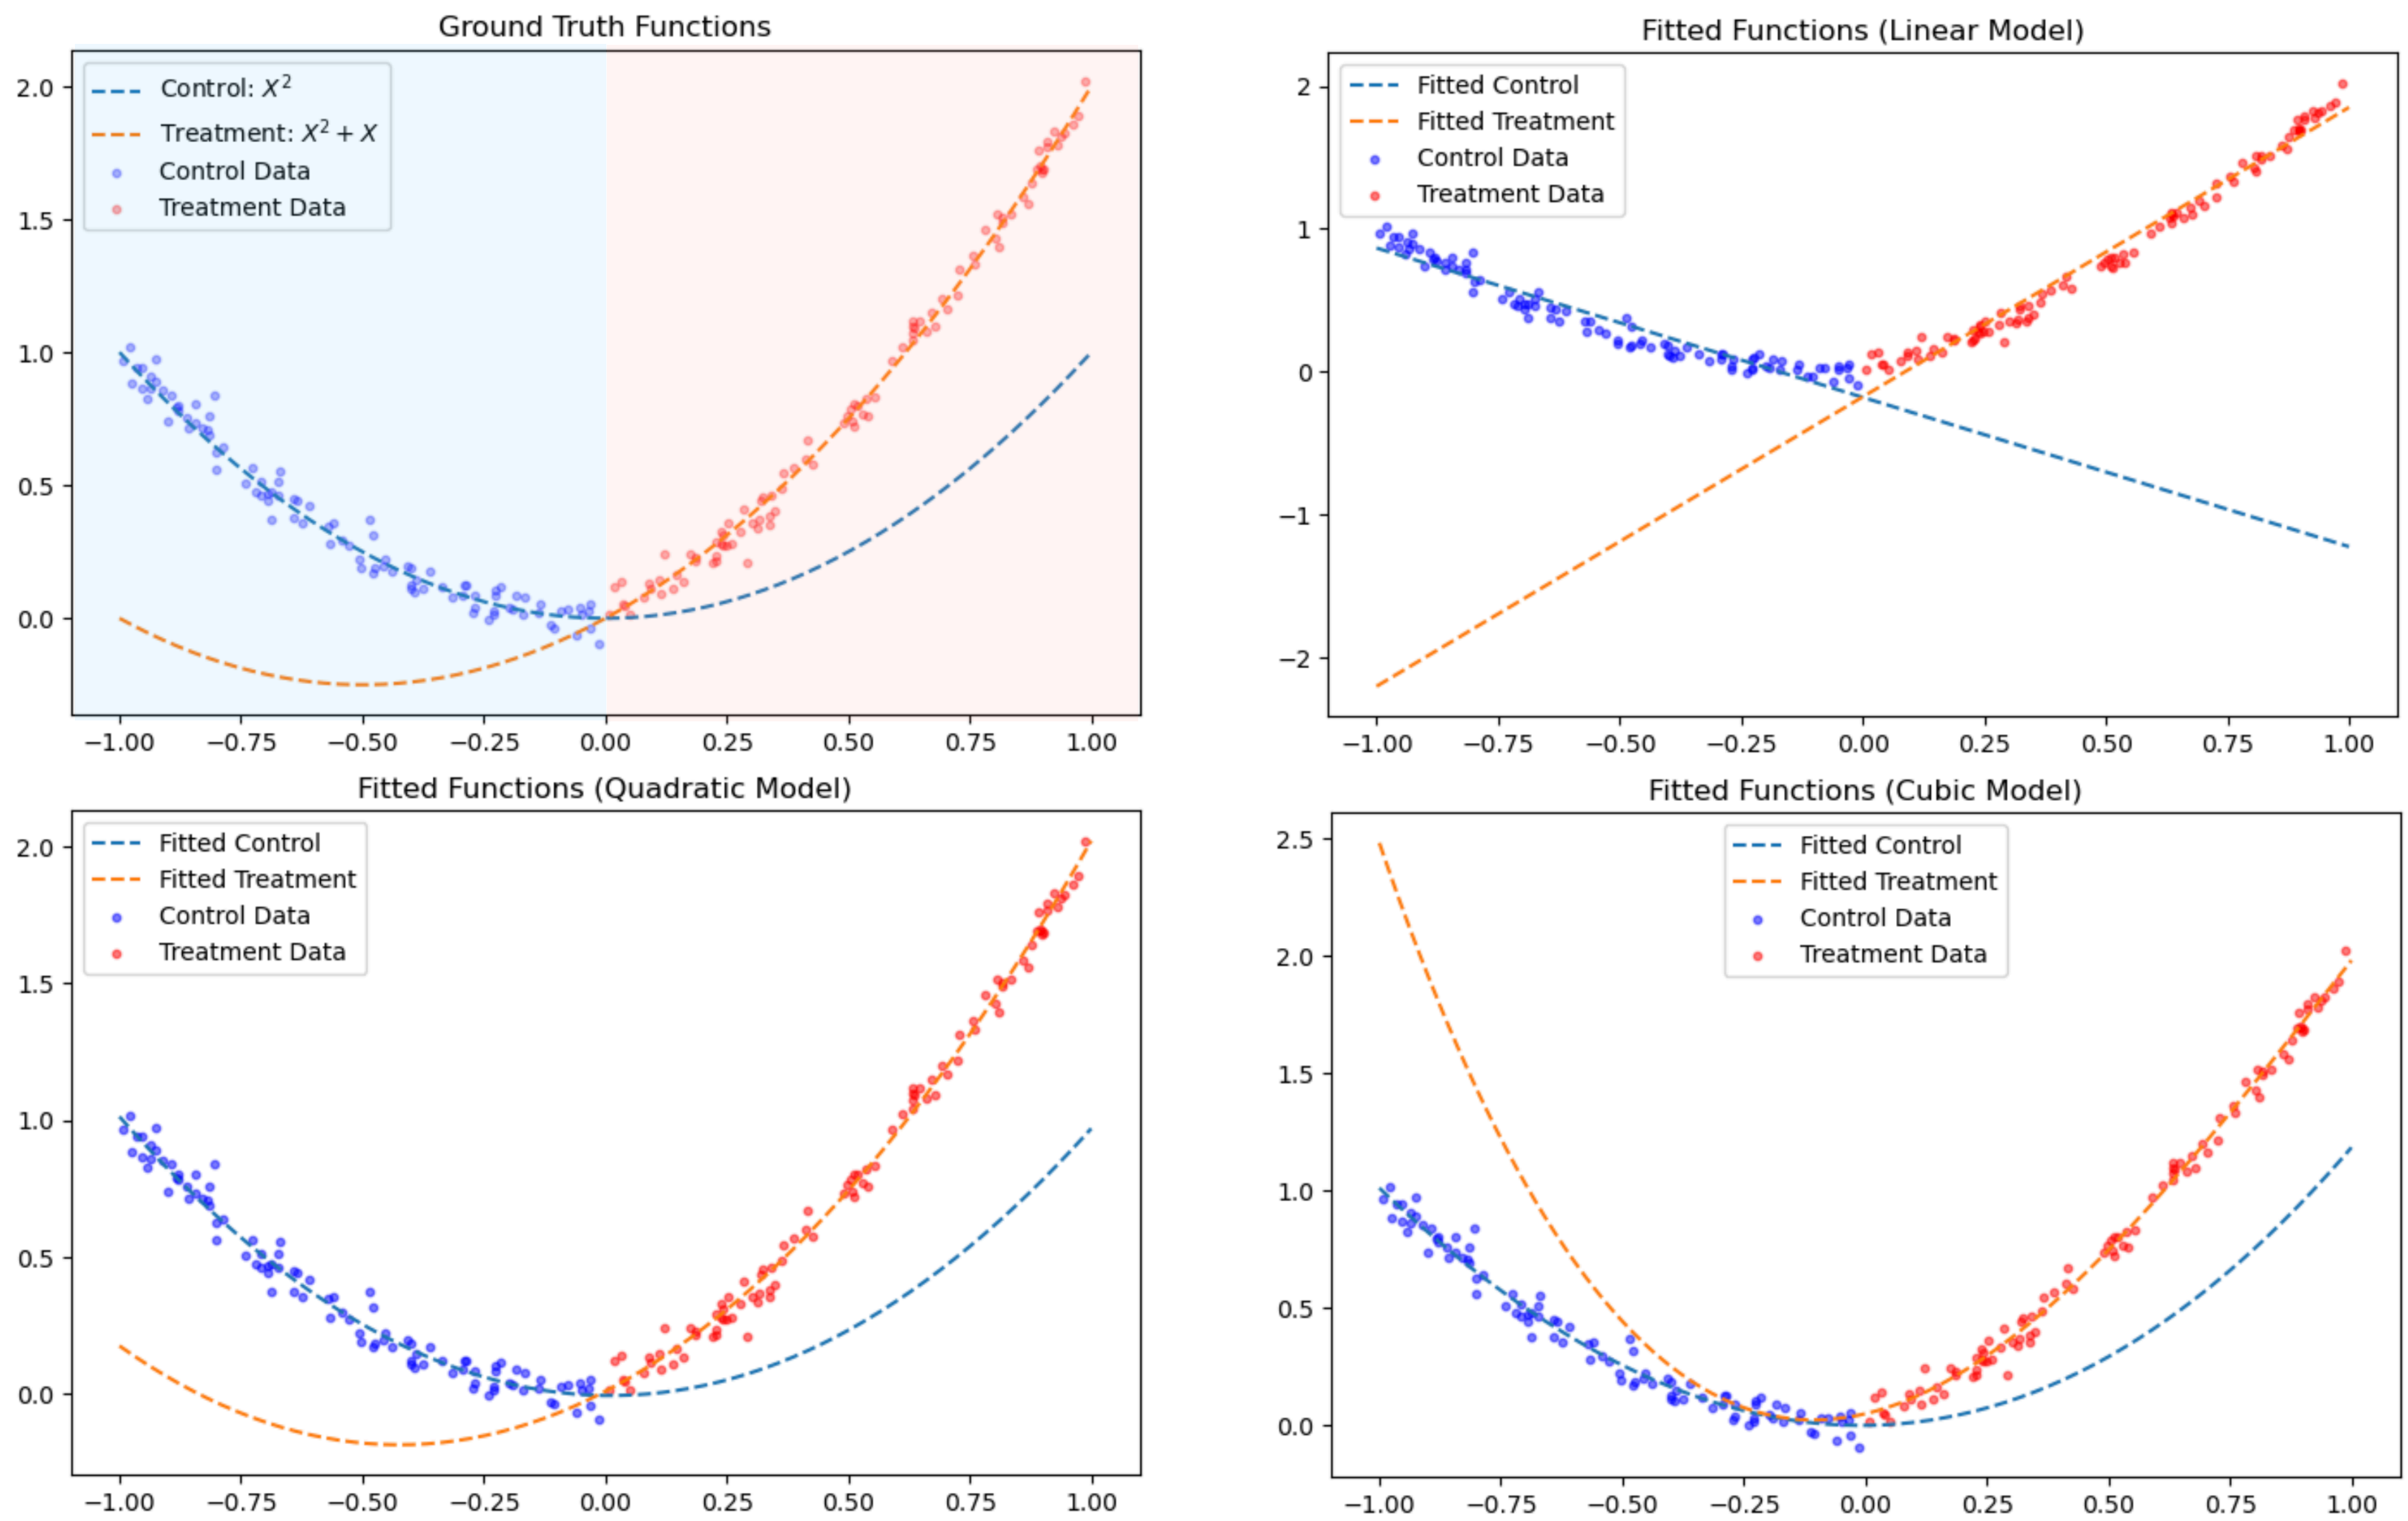

Overlap – Risks of extrapolation

$f_0(X) = X^2 + \epsilon$ $\qquad f_1(X) = X^2 + X + \epsilon$ $\qquad \text{CATE}(X) = X$

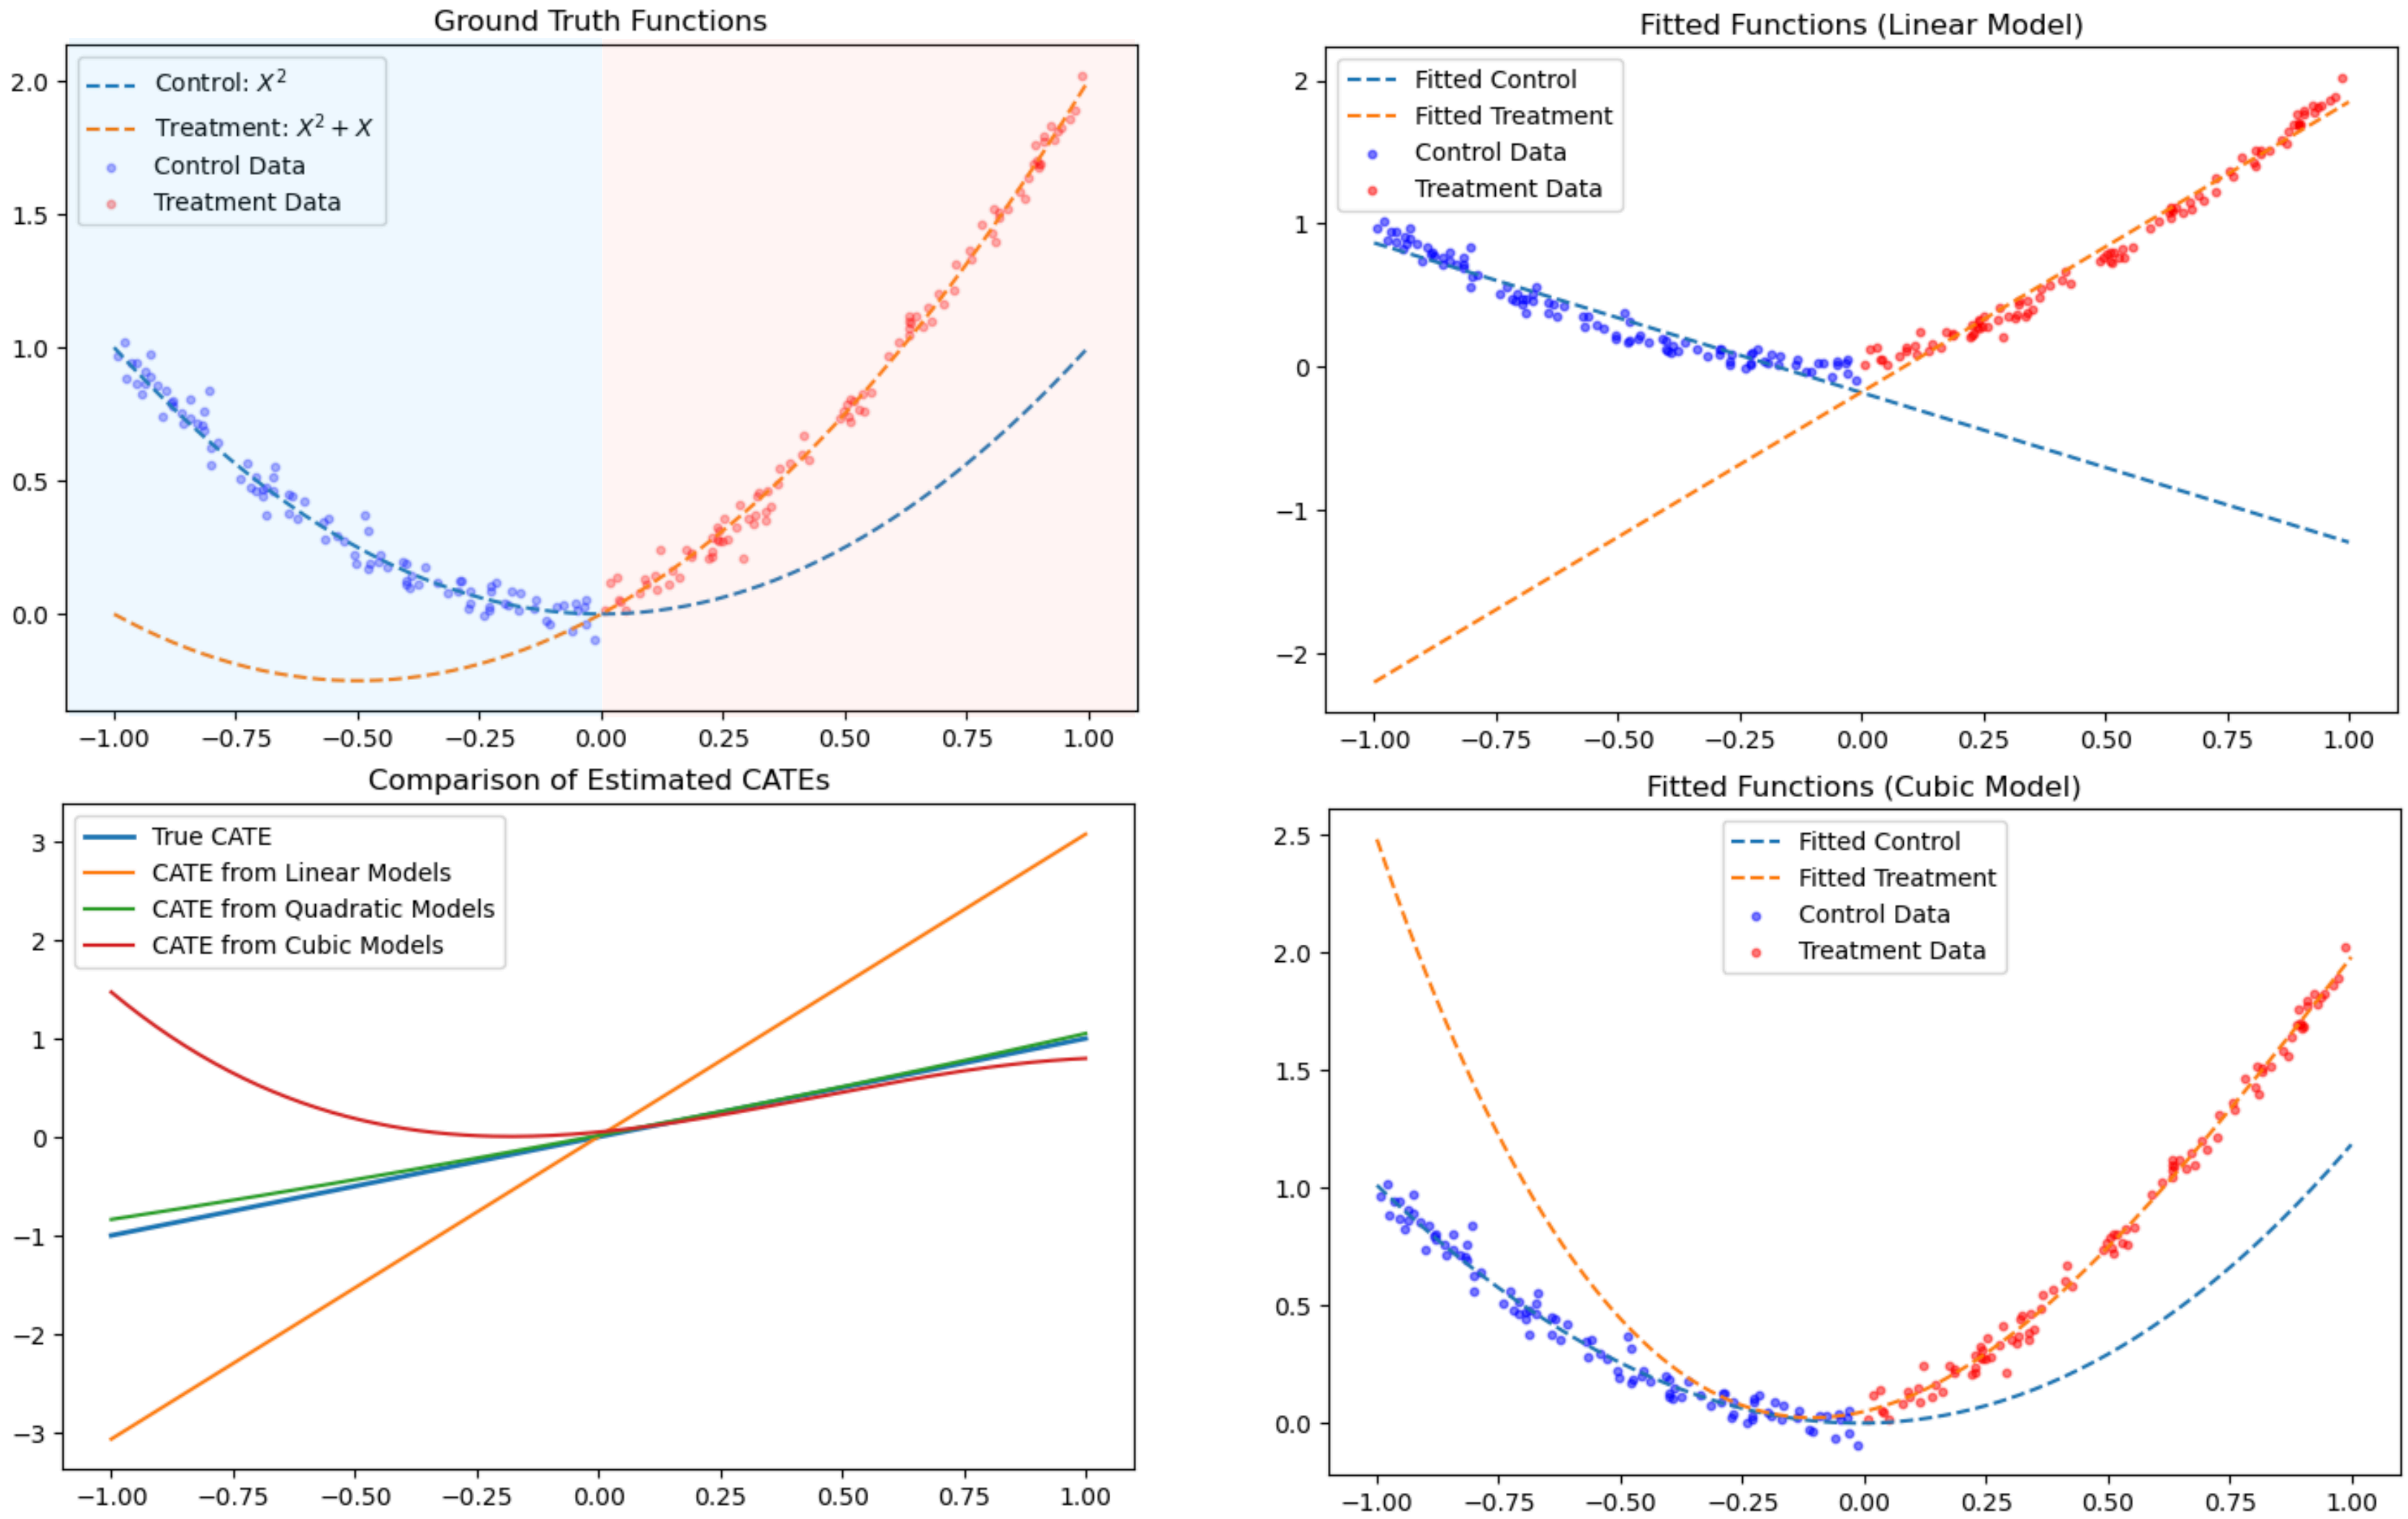

Overlap – Risks of extrapolation

$f_0(X) = X^2 + \epsilon$ $\quad$ $f_1(X) = X^2 + X + \epsilon$ $\quad$ $\text{CATE}(X) = X$

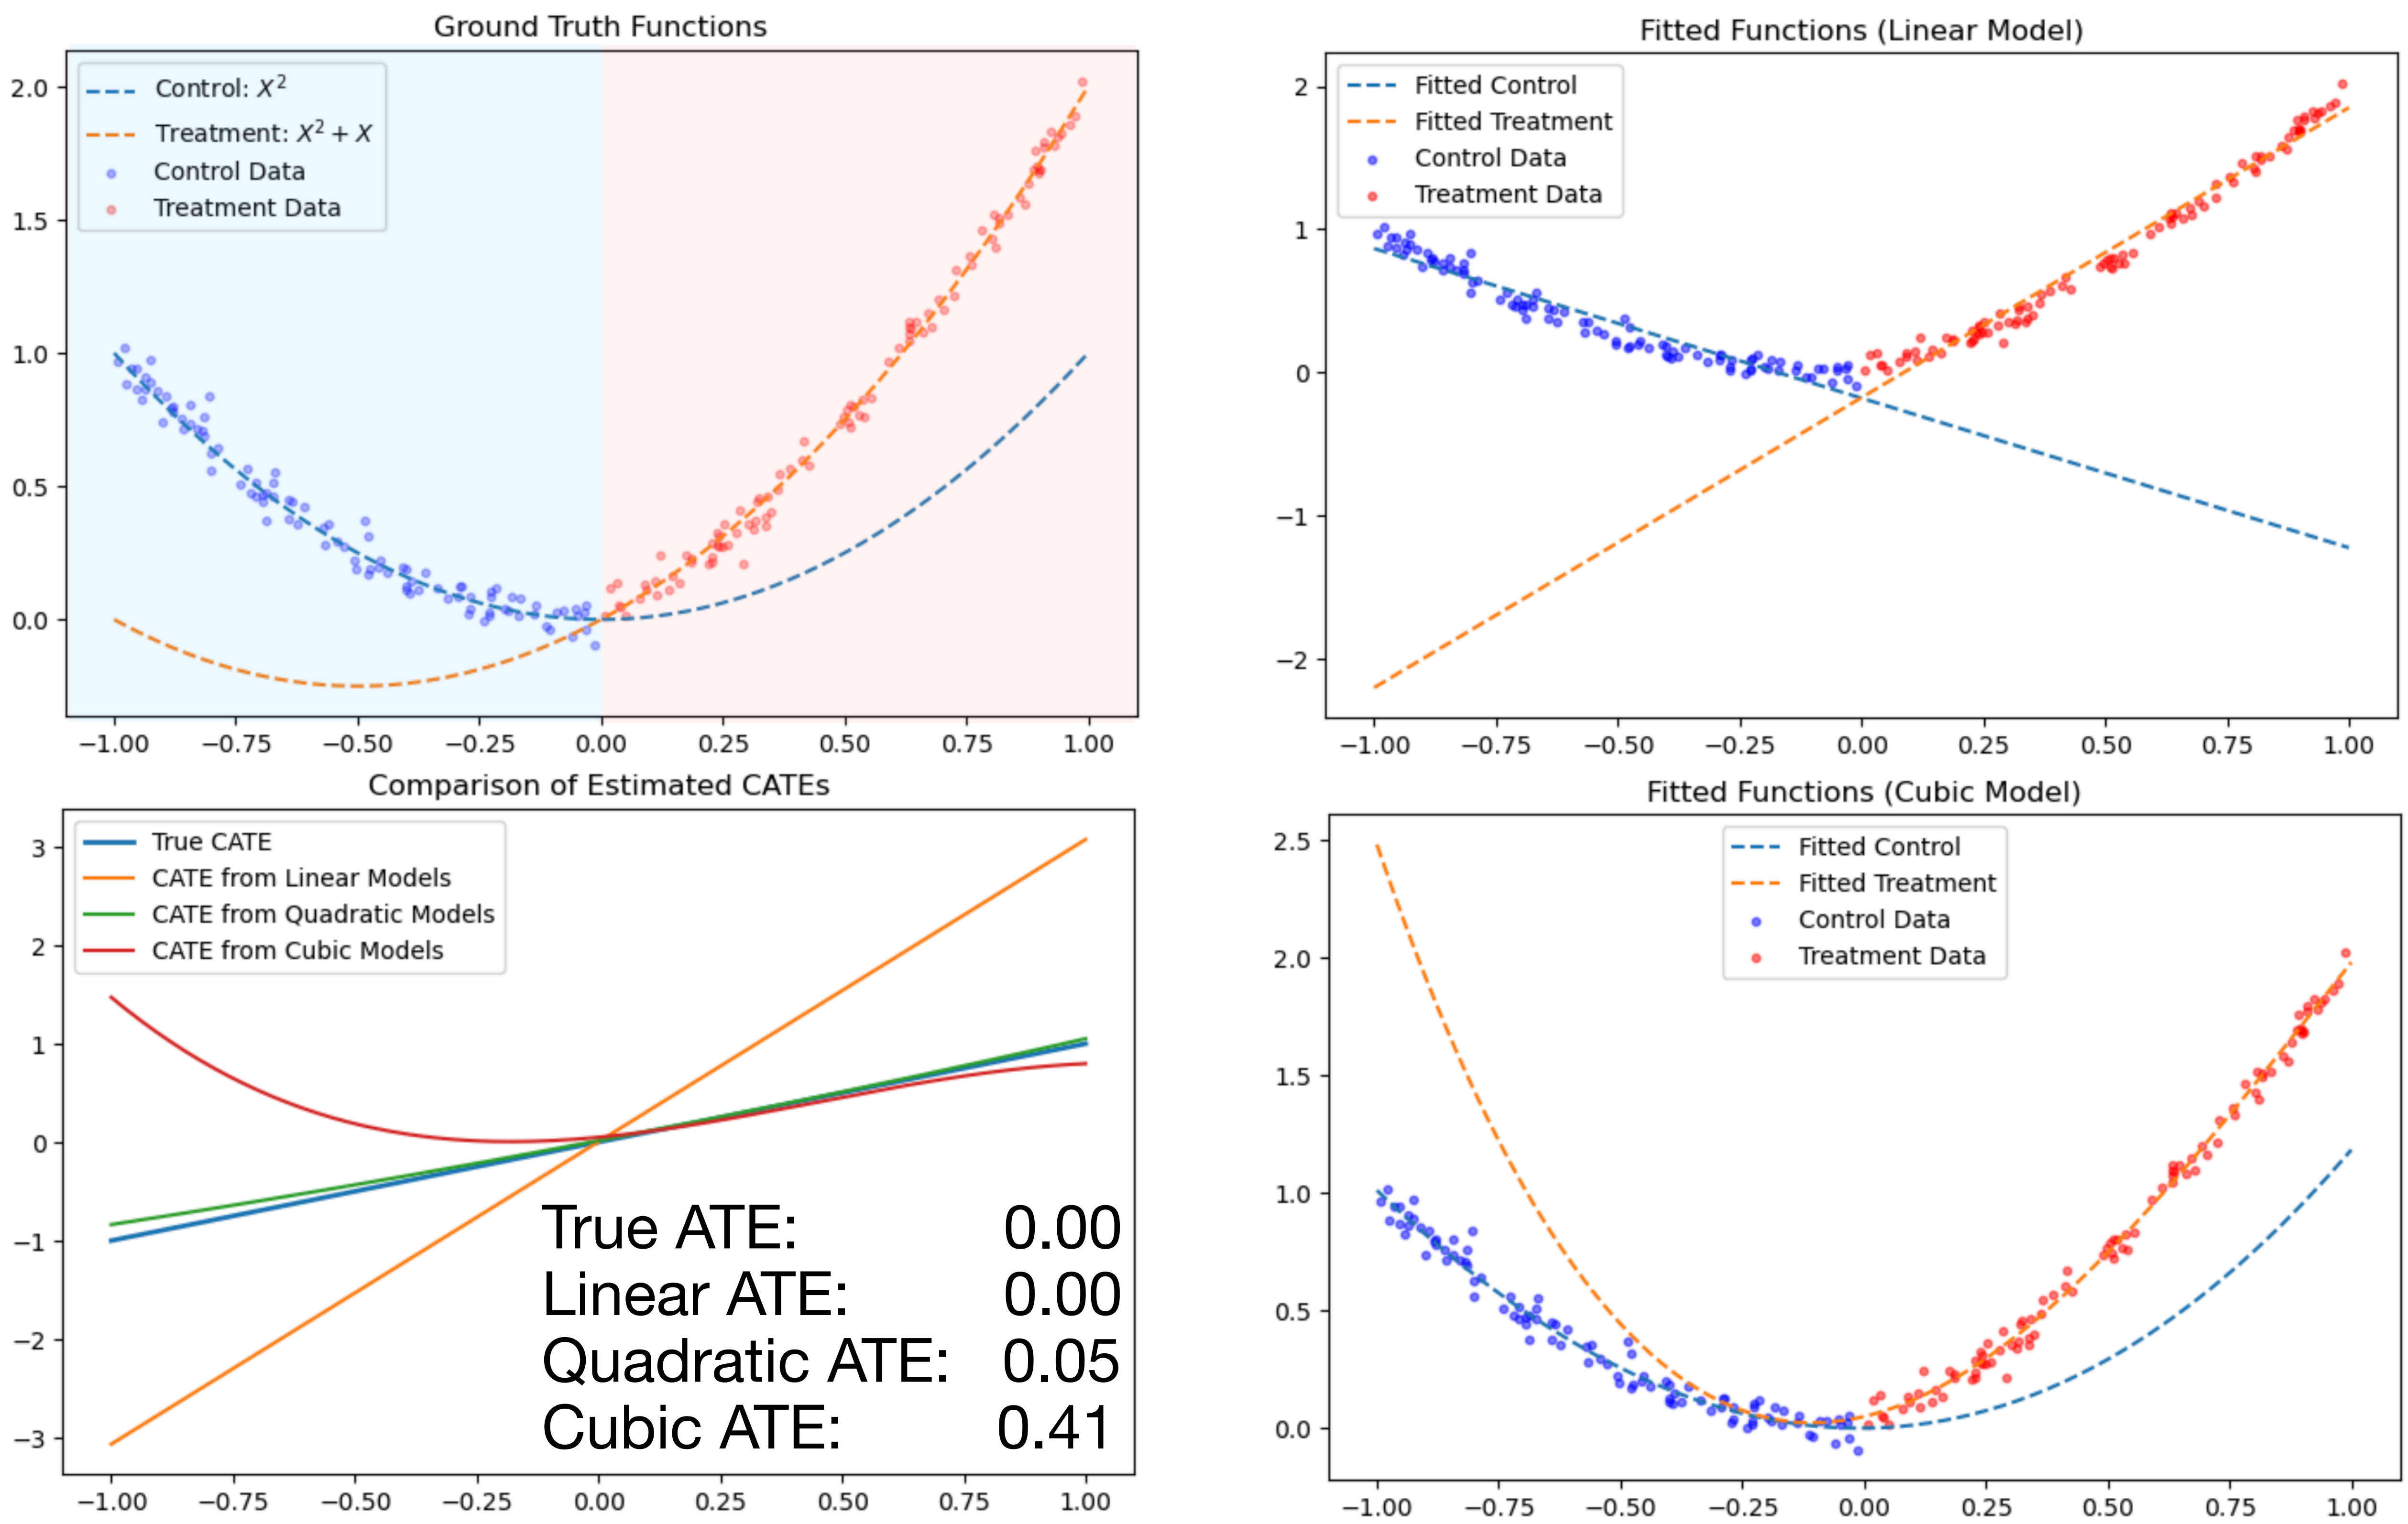

Overlap – Risks of extrapolation

$f_0(X) = X^2 + \epsilon$ $\quad f_1(X) = X^2 + X + \epsilon$ $\quad \text{CATE}(X) = X$

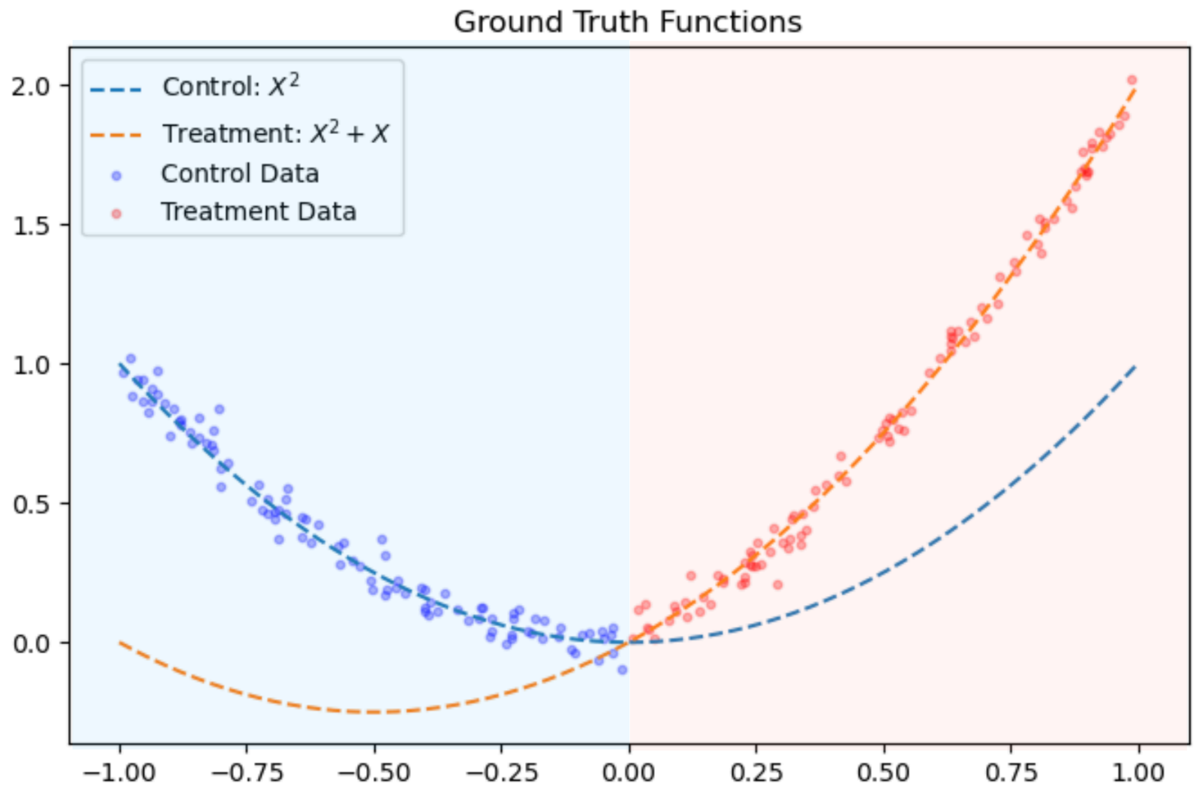

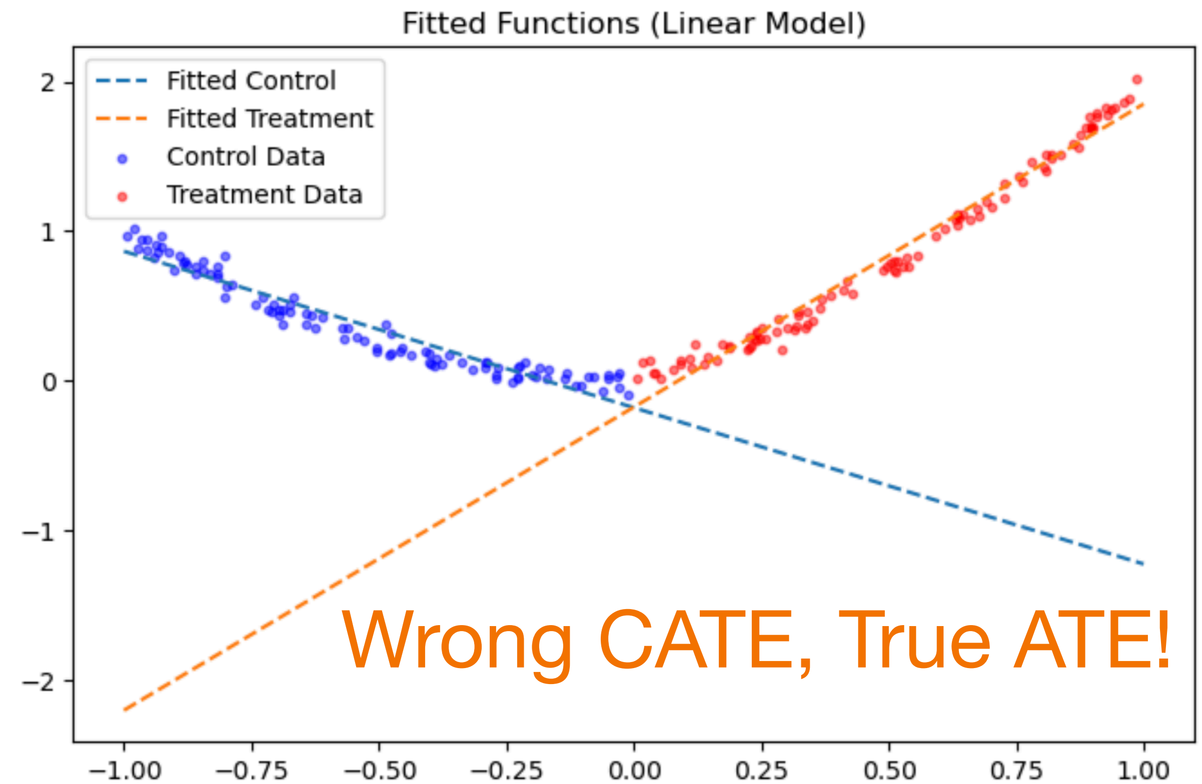

In general, we can assume that the prognostic model, m, is given at the time of the trial. To analyze the trial data, we first use the prognostic model to generate the prognostic score, $M_i = m(X_i)$, for each subject given their baseline covariates, X_i. Then, we estimate the treatment effect using a linear regression adjusted for the empirically centered covariates, prognostic score, and their interactions with the treatment[3]. Letting $Z^\top = [1, \tilde{W}, \tilde{X}^\top, \tilde{M}^\top, \tilde{W}\tilde{X}^\top, \tilde{W}\tilde{M}^\top]$ be the regressors [4], we fit $\mathbb{E}\left[Y_i \middle| Z_i\right] = Z_i^\top \beta$ using ordinary least squares to obtain the fit coefficients, $\hat{\beta}$. Our estimate of the treatment effect is $\hat{\tau} = \hat{\beta}_W$, i.e., the coefficient corresponding to the W term in the regression. This specification is directly based on the "ANCOVA II" estimator analyzed in Yang and Tsiatis [18]. It is well-known that this is a consistent and asymptotically normal estimator of the treatment effect when treatment is randomized, even if the regression is misspecified (i.e., the true relationship is nonlinear)[5] [25, 26, 18].

[Increasing the efficiency of randomized trial estimates via linear adjustment for a prognostic score]

Overlap – IPW estimators

What is the implication of lack of overlap for IPW estimators?

$$A\hat{T}E = \frac{1}{n}\sum_{i \text{ s.t. } t_i=1}\frac{y_i}{\hat{p}(t_i = 1|x_i)} - \frac{1}{n}\sum_{i \text{ s.t. } t_i=0}\frac{y_i}{\hat{p}(t_i = 0|x_i)}$$

Overlap – IPW estimators

What is the implication of lack of overlap for IPW estimators?

$$\hat{ATE} = \frac{1}{n} \sum_{i \text{ s.t. } t_i=1} \frac{y_i}{\hat{p}(t_i = 1|x_i)} - \frac{1}{n} \sum_{i \text{ s.t. } t_i=0} \frac{y_i}{\hat{p}(t_i = 0|x_i)}$$

Vanishing denumerators

Exploding terms

High variance in the estimate

Overlap – IPW estimators

What is the implication of lack of overlap for IPW estimators?

$$\hat{ATE} = \frac{1}{n} \sum_{i \text{ s.t. } t_i=1} \frac{y_i}{\hat{p}(t_i = 1|x_i)} - \frac{1}{n} \sum_{i \text{ s.t. } t_i=0} \frac{y_i}{\hat{p}(t_i = 0|x_i)}$$

Vanishing denumerators

Exploding terms $\implies$ Clip $\hat{p}(t \mid x)$ from above to 0.95 and from below to 0.05

High variance in the estimate

What is the problem with clipping?

Overlap – IPW estimators

What is the implication of lack of overlap for IPW estimators?

$$\hat{ATE} = \frac{1}{n}\sum_{i \text{ s.t. } t_i=1} \frac{y_i}{\hat{p}(t_i = 1|x_i)} - \frac{1}{n}\sum_{i \text{ s.t. } t_i=0} \frac{y_i}{\hat{p}(t_i = 0|x_i)}$$

Vanishing denumerators

Exploding terms $\Longrightarrow$ Clip $\hat{p}(t \mid x)$ from above to 0.95 and from below to 0.05

High variance in the estimate

What is the problem with clipping?

Biases the estimates, especially for “atypical” patients.

No free lunch – Everything we do to make things work has consequences. Understanding what those are is important.

Overlap – How to check?

You will do some of this in the Pset 3

Overlap – How to check?

You will do some of this in the Pset 3

Look at the marginals -> Good sanity check

Might not be enough, why?

Overlap – How to check?

You will do some of this in the Pset 3

Look at the marginals -> Good sanity check

Might not be enough, why?

Joint distributions can still be non-overlapping

Especially problematic for high dim. X

In practice, we might be relying on “extrapolation/generalization” abilities of ML models more than we think.

They are getting *really* good at that (*e.g.*, LLMs) – open research area.