Causal Inference with Latent Treatment, Outcome, and Confounder

PS690 Computational Methods in Social Science

Jiawei Fu
Department of Political Science
Duke University

October 28, 2025

Overview

1. General Framework of Latent outcome
2. Latent Outcome from Text
3. Latent Confounder
4. Proximal Causal Inference with Text Data
5. Latent Treatment

Latent Treatment Effects

- Many interesting and important outcomes in the social sciences are abstract concepts that cannot be directly observed.
- Constructs such as economic inequality, press freedom, political violence, corruption, and post-materialism are just a few examples of latent variables that are thought to be measured imperfectly, despite the sustained efforts of researchers.
- Question: How should we analyze randomized experiments in which the main outcome is measured in multiple ways and each measure contains some degree of error?

Current Practice

- With multiple measurement, current practice are
  1. Principal components analysis (PCA)
  2. Anderson's inverse-covariance weighting (ICW)
  3. Item Response Theory (IRT)
  4. Structural Equation Modeling (SEM)
  5. Seemingly Unrelated Regression (SUR)
- Key problem:
  1. Some of them ignore the latent variable and structure. They are not estimating latent treatment effects. They just reduce dimensionality.
  2. Latent variable has no intrinsic unit or scale. The arbitrary dimension reduction results in a situation where, even though two studies aim to estimate the same latent effects, their results are not comparable.
  3. To be specific, even though the researchers focus on the say latent treatment, because of different measurements in different studies, arbitrary dimension reduction distort the interpretation of latent variable, which results in different scale.

Framework

- Treatment: Z_i
- Potential latent variables: η_i^1
- The latent outcome η_i can be represented as $\eta_i = \mathbb{E}\eta_i^0 + \tau Z_i + \zeta_i$, where ζ_i is the idiosyncratic disturbance:
  $\zeta_i = \eta_i^0 - \mathbb{E}\eta_i^0 + Z_i[(\eta_i^1 - \mathbb{E}\eta_i^1) - (\eta_i^0 - \mathbb{E}\eta_i^0)]$.
- j^{th} outcome measure for unit i be
  $Y_{ij} = \lambda_j \eta_i + \epsilon_{ij} =$
  $\lambda_j[Z_i\eta_i^1 + (1 - Z_i)\eta_i^0] + [Z_i\epsilon_{ij}^1 + (1 - Z_i)\epsilon_{ij}^0]$

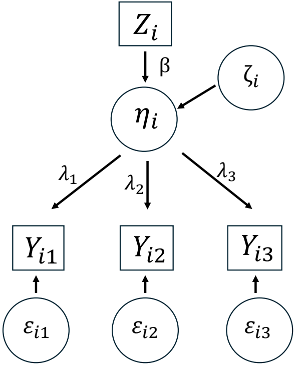

Figure: [Fu and Green, 2025]

Framework

- Individual Latent Treatment Effect (LTE): $\tau_i = \eta_i^1 - \eta_i^0$
- Average LTE for a given set of subjects is $\tau = \frac{1}{n}\sum_{i=1}^{n}\eta_i^1 - \frac{1}{n}\sum_{i=1}^{n}\eta_i^0 = \frac{1}{n}\sum_{i=1}^{n}\tau_i$;
- The super-population average treatment effect is $\mathbb{E}[\eta_i^1 - \eta_i^0]$.

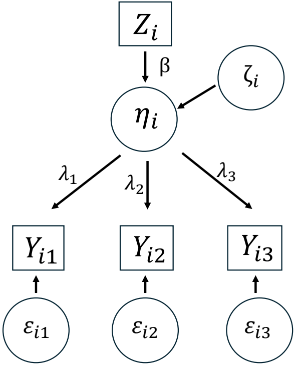

Figure: [Fu and Green, 2025]

Framework

Assumption (Causal Framework for a Latent Outcome Variable)

In the Causal Framework for a Latent Outcome Variable, we assume $\forall j = 1, 2, ..., J$, *and* $i = 1, 2, ..., n$,

A. Valid Measurement: $\lambda_j \neq 0$.

B. Z_i *is randomly assigned:* $Var(Z_i) \neq 0$ *and* $\{\eta_i^0, \eta_i^1\} \perp Z_i$

C. Z_i *is excludable:* $\epsilon_{ij} \perp Z_i$

D. Stable unit treatment value assumption (SUTVA):

- *a. If* $Z_i = Z_i'$, *then* $\eta_i(Z_1, Z_2, ..., Z_i, ..., Z_n) = \eta_i(Z_1, Z_2, ..., Z_i', ..., Z_n)$
- *b. If* $Z_i = z$, *then* $\eta_i = \eta_i(z)$, $\forall i$ *and* $z \in Z$.

Framework

- Key Identification assumption: $\lambda_1 = 1$, so that we can interpret η using the same metric as Y_1.
- For example, if η_i represents a latent distance, and Y_{i1} is scaled in terms of kilometers while Y_{2i} is scaled in terms of miles, setting $\lambda_1 = 1$ means that effects on η_i are scaled in terms of kilometers.
- Recall, standardization imposes scaling parameters that are sample-specific: Consider two very large studies in which the intervention exerts the exactly same average treatment effect on a latent outcome measured by the same Y_{ij} indicators.
- Standardization would produce two quite different estimates of the ALTE if the measurement error variances were much larger in the first study than the second.
- Standardization can lead a researcher to mistakenly conclude that two estimated ALTEs are different when they are in fact the same.

Identification

- Therefore, if we knew λ_j, we could rescale the observed measures, for example, by $\frac{Y_{ij}}{\lambda_j} = \eta_i + \frac{1}{\lambda_j}\epsilon_{ij}$, to approximate the latent variable η_i using the same units as Y_{i1}.
- The question is how to identify λ_j.
- Recall that $Y_{i2} = \lambda_2(\eta_i^0 + Z_i\tau_i) + \epsilon_{i2}$. This equality implies the following relationship between Y_{i2} and Y_{i1}

$$Y_{i2} = \lambda_2 Y_{i1} + (\epsilon_{i2} - \lambda_2\epsilon_{i1})$$

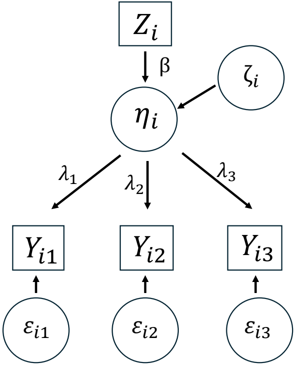

Figure: [Fu and Green, 2025]

Identification

$$Y_{i2} = \lambda_2 Y_{i1} + (\epsilon_{i2} - \lambda_2 \epsilon_{i1})$$

- Note that Y_{i1} is endogenous here.
- It turns out that Z_i can be a valid Instrumental variable under excludability assumption.

$$Cov(Z_i, Y_{i2}) = \lambda_2 Cov(Z_i, Y_{i1}) + Cov(Z_i, \epsilon_{i2} - \lambda_2 \epsilon_{i1})$$

- Under same logic, we can observe that Y_{i3} is also a valid IV.
- Therefore, all λ are (over-)identified.

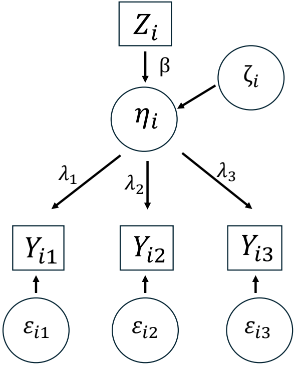

Figure: [Fu and Green, 2025]

Estimation

- With identified λ, we can construct transformed outcomes let $\tilde{Y}_{ij} = \frac{1}{\lambda_j} Y_{ij}$, and $\tilde{Y}_i = \sum_{j=1}^{J} \omega_j \tilde{Y}_{ij}$, where ω_j is the weight and $\sum_{j=1}^{k} \omega_j = 1$.
- Then, ALTE can be estimated by the difference-in-means estimator:

$$
\begin{aligned}
\hat{\tau} &= \frac{1}{n_1} \sum_{i=1}^{n} Z_i \tilde{Y}_i - \frac{1}{n_0} \sum_{i=1}^{n} (1 - Z_i) \tilde{Y}_i \\
&= \frac{1}{n_1} \sum_{i=1}^{n} Z_i [\sum_{j=1}^{k} \omega_j \tilde{Y}_{ij}] - \frac{1}{n_0} \sum_{i=1}^{n} (1 - Z_i) [\sum_{j=1}^{k} \omega_j \tilde{Y}_{ij}]
\end{aligned}
$$

- We can also calculate the optimal weights ω by minimizing the variance.
- To account for uncertainty of estimated $\hat{\lambda}$, similar to IPW estimator, we need to use GMM.

Estimation

- The proposed method has largest power than current methods.

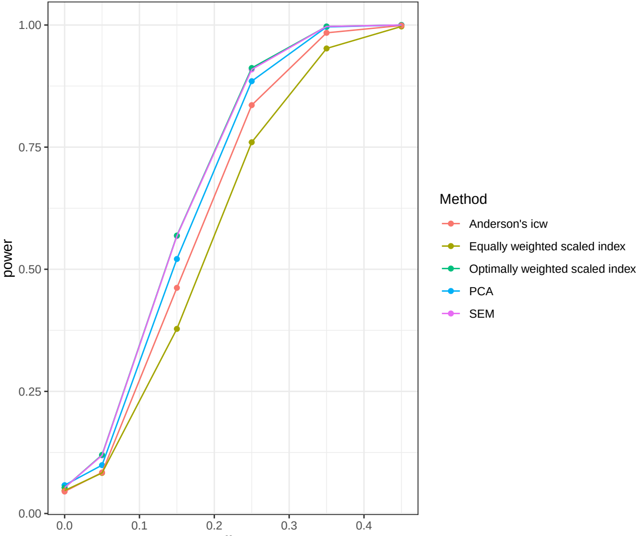

Trade-off

- Researchers could improve the precision of their experiments by investing in additional outcome measures.
- This is analogous to adding more control variables or increasing the sample size.
- We can apply this framework to a more general budget allocation problem.
- Suppose that researchers have a budget $B > 0$ and the cost of adding one more measure is c_1 and the cost of adding one more observation is c_2. Consider the optimal allocation of budget for the sample size n and the number of items J.
- It turns out that measurement reliability affects the marginal gain of measurements.

Trade-off

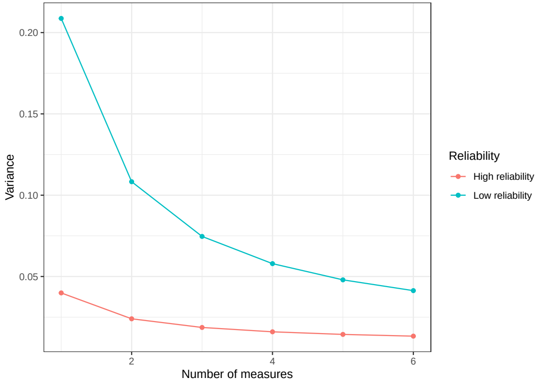

- For high reliability case, the optimal solution is to add one more measure and recruit 400 more respondents.
- For low reliability case, the optimal solution is to add two more measures and recruit 300 more respondents

Text as outcome

- Now, the outcomes of interest are represented within a text document Y.
- Documents are high-dimensional, complicated, and sparse; hence, text is typically not usable for social science inference in its raw form.
- In this case, we need a g function to define the low-dimensional categories of the outcome.
- Then we can apply the traditional difference-in-means estimator.
- For example, authors study how knowledge about an individual's criminal history affects respondent's preference for punishment and deportation.
- We can use STM to discover topics in the survey open-ended response, and use binary coding for each topic as an outcome.

Text as outcome

	Label	Highest probability words
Topic 1	Limited punishment with help to stay in country, complaints about immigration system	legal, way, immigr, danger, peopl, allow, come, countri, can, enter
Topic 2	Deport	deport, think, prison, crime, alreadi, imprison, illeg, sinc, serv, time
Topic 3	Deport because of money	just, send, back, countri, jail, come, prison, let, harm, money
Topic 4	Depends on the circumstances	first, countri, time, came, jail, man, think, reason, govern, put
Topic 5	More information needed	state, unit, prison, crime, immigr, illeg, take, crimin, simpli, put
Topic 6	Crime, small amount of jail time, then deportation	enter, countri, illeg, person, jail, deport, time, proper, imprison, determin
Topic 7	Punish to full extent of the law	crime, violent, person, law, convict, commit, deport, illeg, punish, offend
Topic 8	Allow to stay, no prison, rehabilitate, probably another explanation	dont, crimin, think, tri, hes, offens, better, case, know, make
Topic 9	No prison, deportation	deport, prison, will, person, countri, man, illeg, serv, time, sentenc
Topic 10	Should be sent back	sent, back, countri, prison, home, think, pay, origin, illeg, time
Topic 11	Repeat offender, danger to society	believ, countri, violat, offend, person, law, deport, prison, citizen, individu

Text as outcome

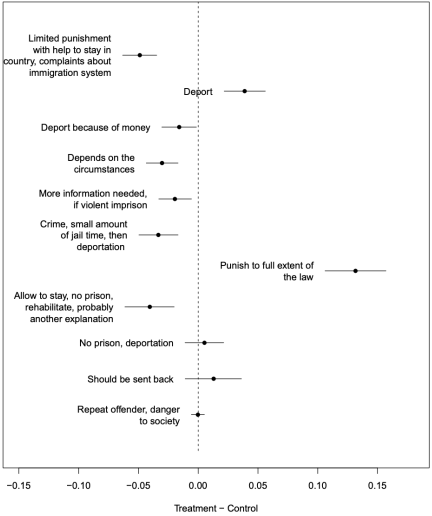

Discovering g

- While g is necessary to make causal inference, it is rarely known exactly from a theory or prior research.
- Instead, g is typically developed through iteration between coding rules and the documents to be coded.
- There are three strategies for learning g from the data:
  1. We could read a sample of text.
  2. We could use supervised learning, to infer g from hand-coded or otherwise labeled documents.
  3. We could use unsupervised learning techniques to discover a low-dimensional representation.

The problem of causal inference with g

- So far, we just assume that we already have a g in hand in the previous identification of ATE.
- However, g is often discovered by interacting with some of the data. We denote the set of documents considered in development of g as J and write g_J to indicate the dependence of g on the documents.
- Sometimes, the discovery of g depends on the randomization.
- Moreover, the set of documents used to develop g, J, overlaps with the set of documents used in estimation, which we will call I.
- Then, researchers might overfit: discover effects that are present in a particular sample but not in the population.

A stylized experiment with text-based outcomes

Table 1. A stylized experiment with text-based outcomes. (**A**) shows the potential outcomes for each unit under each treatment assignment. Treated units talk about candidate morals and polarization and control units talk about taxes and immigration. In (**B**), *T* denotes a different treatment assignment vector where two of four units are treated. *Y* denotes text-based observed outcomes under each treatment assignment. Mo, Im, Tx, and Po stand for candidate morals, immigration, taxes, and polarization, respectively.

(A) Text-based potential outcomes

	Potential outcome under treatment	**Potential outcome under control**
Person 1	Candidate morals	Taxes
Person 2	Candidate morals	Taxes
Person 3	Polarization	Immigration
Person 4	Polarization	Immigration

(B) Text-based observed outcomes under six different treatment assignments

	T	Y	T	Y	T	Y	T	Y	T	Y	T	Y
Person 1	1	Mo	1	Mo	1	Mo	0	Tx	0	Tx	0	Tx
Person 2	1	Mo	0	Tx	0	Tx	1	Mo	1	Mo	0	Tx
Person 3	0	Im	1	Po	0	Im	1	Po	0	Im	1	Po
Person 4	0	Im	0	Im	1	Po	0	Im	1	Po	1	Po
Number of categories	2		4		4		4		4		2	

Figure: From [Egami et al., 2022]

The problem of causal inference with g

- The most straightforward approach is to define g before looking at the documents.
- The preferred procedure is to explicitly separate the creation of g and the estimation of treatment effects.
- We randomly divide our data into a two sets. Specifically, we randomly create a set of units in one set denoted by the indices J and a nonoverlapping set denoted by the indices I.
- We use only the J set to estimate the g_J function and then discard it. We then use the I set exclusively to estimate the causal effect.

Text as Confounder

- How to use text data to recover latent confounders?
- One approach requires pre-specified confounders, which can be learned through g function.
- To fix idea, consider a study on the effects of experiencing government censorship on Chinese social media users.
- We ask (1) are Chinese social media users who have a post censored more likely to be censored in subsequent posts? (2) does censorship decrease the number of future posts by a user?
- Researchers may think that only topics of the texts and individual words are confounders.
- Then, a useful g function is one that retains information about both of these channels and eliminates other information about the text.

Text as Confounder

- Problem: it is not that easy to identify all confounders.
- Moreover, given the g function captures all confounders, since NLP methods are still far from perfectly accurate, how can we mitigate error that arises from approximating confounding variables?
- In the big data era, people find that proximal causal inference is more attractive.
- Indeed, we can apply similar idea to text data.

Review of Proximal Causal Inference

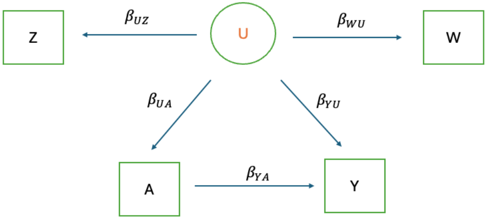

- A is the treatment and Y is the outcome. We assume STUVA holds.
- Negative control outcome (NCO): $W \perp\!\!\!\perp A | U$
- Negative control exposure (NCE): $Z \perp\!\!\!\perp Y | (U, A)$ and $Z \perp\!\!\!\perp W | (U, A)$
- NCO and NCE should associated with U; variation in U can be recovered from variation in Z and W.

Identification with NCO

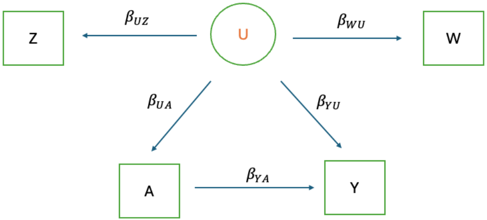

- Suppose regress Y on A, the effect is confounded by U: $\beta_{YA} + \beta_{YU}\beta_{UA}$.
- Now, regress NCO W on A, the coefficient is $\beta_{WU}\beta_{UA}$, which reflects the association due to U.
- If we assume the same association between $U - Y$ and $U - W$: $\beta_{YU} = \beta_{WU}$, then the bias $\beta_{YU}\beta_{UA} = \beta_{WU}\beta_{UA}$ is identified and so is β_{YA}.

Identification with NCE

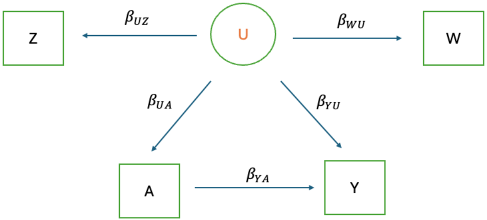

- Again, regress Y on A, we have baised effect: $\beta_{YA} + \beta_{YU}\beta_{UA}$.
- Now, regress Y on NCE Z, the coefficient is $\beta_{YU}\beta_{UZ}$, which reflects the association due to U.
- If we assume the same association between $U - A$ and $U - Z$: $\beta_{UA} = \beta_{UZ}$, then the bias $\beta_{YU}\beta_{UA} = \beta_{YU}\beta_{UZ}$ is identified and so is β_{YA}.

Identification with Double Negative Controls

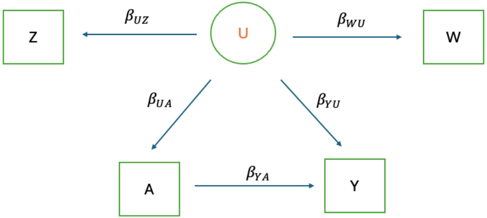

- Suppose the previous equal association assumptions do not hold.
- Regress NCO W on A, the coefficient $\beta_{WU}\beta_{UA}$ reflects confounding bias up to a scale which is $\frac{\beta_{YU}}{\beta_{WU}}$.
- This ratio can be identified using NCE: $\frac{Y \sim Z}{W \sim Z} = \frac{\beta_{YU}\beta_{UZ}}{\beta_{WU}\beta_{UZ}}$.
- Therefore, β_{YA} is identified.

Proximal Causal Inference

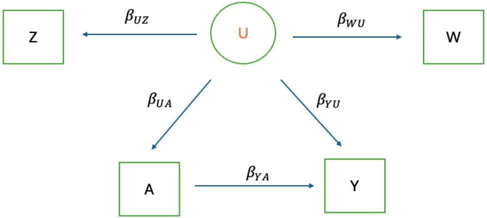

- More generally, we assume exist a function $h(W, A)$ that captures the relationship between $U - Y$ and $U - W$: $\mathbb{E}[Y|A = a, U] = \mathbb{E}[h(W, a)|A = a, U]$.
- It means that the confounding effect of U on Y at $A = a$ is equal to the confounding effect of U on $h(W, a)$, a transformation of W.
- Then, ATE can be identified: $\mathbb{E}[Y(a)] = \mathbb{E}[h(W, A = a)] \ \forall a$.
- Why? Take expectation over U:
  $LHS = \mathbb{E}[\mathbb{E}[Y|A = a, U]] = \mathbb{E}[\mathbb{E}[Y(a)|U]] = \mathbb{E}[Y(a)]$. And
  $RHS = \mathbb{E}[\mathbb{E}[h(W, A = a)]|A = a, U] = \mathbb{E}[\mathbb{E}[h(W, a)|U]] = \mathbb{E}[h(W, a)]$.

Proximal Causal Inference

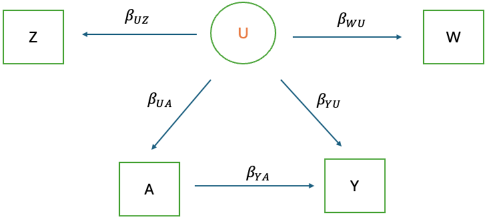

- How to identify $h(W, A)$?
- Given $\mathbb{E}[Y|A, U] = \mathbb{E}[h(W, A)|A, U]$, taking expectation of U with respect to $p(U|A, Z)$ on both sides, we obtain $\mathbb{E}[Y|A, Z] = \mathbb{E}[h(W, A)|A, Z]$.
- It shows that $h(W, A)$ also captures the relationship between $Z - Y$ and $Z - A$.
- Therefore, $h(W, A)$ is called outcome confounding bridge function.

Proximal Causal Inference

- Note that $\mathbb{E}[Y - h(W, A = a)|A = a, Z] = 0$.
- To get unique h, we need some more assumptions, for example $\mathbb{E}[g(W)|A = a, Z = z] = 0 \; \forall z$, then $g(W) = 0$.
- How to estimate $h(W, A)$?
- We can try parametric function first, and use GMM to estimate all parameters.
- We can also try more flexible approaches such as semi-parametric or nonparametric methods.

Proximal Causal Inference

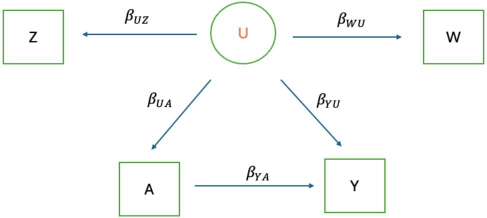

- If assume a system of linear functions

$$\mathbb{E}[Y|A,U] = \beta_{Y0} + \beta_{YA}A + \beta_{YU}U$$
$$\mathbb{E}[W|U] = \beta_{W0} + \beta_{WU}U$$

- Let $\tilde{U} = \frac{W - \beta_{W0}}{\beta_{WU}}$. $\mathbb{E}[\tilde{U}|U] = U$. Then $\mathbb{E}[Y|A,Z] = \beta_{Y0} + \beta_{YA}A + \beta_{YU}\mathbb{E}[\tilde{U}|A,Z]$.
- Thus, a natural function of h is $\theta_0 + \theta_A A + \theta_W W$. Here, we can regress W on A, Z; get predicted value $\hat{W}$ and regress Y on A adjusting $\hat{W}$.

Proximal Causal Inference with Text Data

- Now, we want to learn W and Z from text data T.
- There are several things to be noted.
- First, text data should be pre-treatment: otherwise, Z and W are not valid negative controls. See figure below, where C is observed covariates.
- Second, use separate texts to learn Z and W. Otherwise, it is possible that they W and Z are not independent conditioning on C and U due to same text.

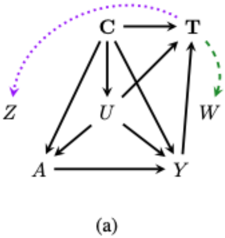

Figure: [Chen et al., 2024]

Text as Treatment

- Follow [Fong and Grimmer, 2023], consider in a study, the treatment is a collection of text X.
- Sometimes, the real interests are some latent treatments within the text: $Z_i = g(X_i) \in \{0, 1\}$, which denotes the presence or absence of the latent treatment in any given document.
- The above function g is called codebook function, which compresses high-dimensional text to a low-dimensional measure used for the treatment.
- The analyst might define g by hand-coding, automatically from the text, by looking for the presence or absence of the word "lawyer," or by a group of words or phrases that convey that someone has a legal background, such as "JD," "attorney," and "law school."

Setup

- Of course, there are often more relevant features in a text than a latent treatment.
- For example, campaign advertisements vary in not only whether they are negative, but also in whether they focus on a candidate's policy positions or character, whether they are light-hearted or ominous, whether they are overt or subtle, and many other potential unmeasured treatments.
- Function $h(X_i)$ to denote other unmeasured latent treatments; define $B_i := h(X_i)$

Setup

- We assume that if two documents have the same measured and unmeasured treatments, then respondents respond to them in the same way, on average.

Assumption

There exists some function h such that if $g(x) = g(x')$ and $h(x) = h(x')$, then $\mathbb{E}[Y_i(X_i = x)] = \mathbb{E}[Y_i(X_i = x')]$*. Additionally,* $0 < \mathbb{P}[Z_i = 1 | B_i = b] < 1$ *for all* $b \in B$*.*

- The first part of the Assumption can also be understood as an independence assumption. Given the measured and unmeasured treatments within a text, the text and the outcome are independent.
- Then, the potential outcome is $Y_i(X_i) = Y_i(g(X_i), h(X_i)) = Y_i(Z_i, B_i)$
- The ATE is

$$ATE = \sum_{b \in B} \{\mathbb{E}[Y_i(Z_i = 1, B_i = b) - \mathbb{E}[Y_i(Z_i = 0, B_i = b)]\} \mathbb{P}(B_i = b)$$

Setup

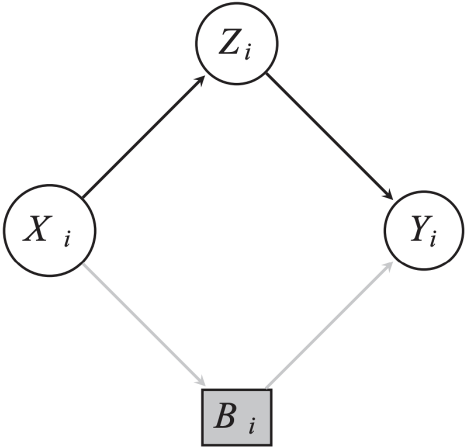

Note: The text, $\boldsymbol{X}_i$, causes both the latent treatment of interest, Z_i, and the unmeasured latent treatments, $\boldsymbol{B}_i$. These latent treatments, in turn, cause the outcome, Y_i.

Assumptions

- If we could directly randomize the latent treatment of interest, then identification and consistent estimation of the ATE is straightforward even if we never directly control for the unmeasured treatments.
- Unfortunately, it is impossible to manipulate the treatment without also manipulating the text. Therefore, we need some identification assumptions
- The first assumption is that an individual's response depends on only the document they are assigned.

Assumption

For all individuals i and any X, X' such that $X_i = X_i'$, $Y_i(X) = Y_i(X')$

Assumptions

- Next, texts are either randomly assigned in an experimental setting or the document assignment is independent of the potential outcomes.

Assumption

For all individuals i, $Y_i(x) \perp\!\!\!\perp X_i$ and $\mathbb{P}[X_i = x] > 0$ for all $x \in X$.

- The above assumption is about who reads a text, not about the contents of a particular text. The assumption is violated if individuals select the particular text they read.

Assumptions

- The key assumption is that

Assumption

At least one of the following is true:

1. *The measured and unmeasured latent treatments are independent:* $\mathbb{P}[Z_i = z, B_i = b] = \mathbb{P}(Z_i = z)\mathbb{P}(B_i = b)$
2. *The unmeasured treatments are unrelated to the outcome:* $\mathbb{E}[Y_i(Z_i = z, B_i = b)] = \mathbb{E}[Y_i(Z_i = z, B_i = b']$ *for all* $z \in \{0, 1\}$ *and all* $b, b' \in B$.

- The conditions in Assumption are analogous to the conditions required to avoid omitted-variable bias in observational research.
- The first condition is that the measured treatments are independent of the unmeasured treatments. If Z_i and B_i are independent, then the distribution of B_i is identical.
- The second condition is that the unmeasured treatments have no effect on the outcome, on average.

Assumptions

- Then, ATE can be identified, and estimated by

$$\mathbb{E}_n[Y_i(X_i)|g(X_i) = 1] - \mathbb{E}_n[Y_i(X_i)|g(X_i) = 0]$$

- The proof is straightforward: WTS that

$$\begin{aligned}
&\mathbb{E}[Y_i(X_i)|g(X_i) = 1] - \mathbb{E}[Y_i(X_i)|g(X_i) = 0] \\
= &\sum_{b \in B} \{\mathbb{E}[Y_i(Z_i = 1, B_i = b) - \mathbb{E}[Y_i(Z_i = 0, B_i = b)]\}\mathbb{P}(B_i = b)
\end{aligned}$$

- The LHS is equal to

$$\begin{aligned}
&\sum_{b \in B} \mathbb{E}[Y_i(X_i)|g(X_i) = 1, h(X_i) = b]\mathbb{P}[h(X_i) = b|g(X_i) = 1] - \\
&\sum_{b \in B} \mathbb{E}[Y_i(X_i)|g(X_i) = 0, h(X_i) = b]\mathbb{P}[h(X_i) = b|g(X_i) = 0]
\end{aligned}$$

Assumptions

- Then, LHS=RHS if

$$\sum_{b \in B} \mathbb{E}[Y_i(X_i)|g(X_i) = 1, h(X_i) = b][\mathbb{P}[B_i = b] - \mathbb{P}[B_i = b|Z_i = 1]] =$$
$$\sum_{b \in B} \mathbb{E}[Y_i(X_i)|g(X_i) = 0, h(X_i) = b][\mathbb{P}[B_i = b] - \mathbb{P}[B_i = b|Z_i = 0]]$$

- If measured and unmeasured latent treatments are independent, then $\mathbb{P}[B_i = b] - \mathbb{P}[B_i = b|Z_i = z] = 0$.
- If unmeasured treatments are unrelated to the outcome, then we can cancel out $\mathbb{E}_n[Y_i(X_i)|g(X_i) = 1, h(X_i) = b]$, and note that $\sum_{b \in B} \mathbb{P}[B_i = b] - \mathbb{P}[B_i = b|Z_i = z] = 0$.

Assumptions

- How to satisfy the previous identification assumption?
- Text-based research, however, has an advantage over researchers grappling with other kinds of omitted variable bias: Every possible confounder is contained within the text.
- In other words, if a researcher can conceive of an unmeasured latent treatment, then the researcher can measure and adjust for it.
- But I still think it's not that easy to implement in practice. More research is needed in this area.

References

Chen, J., Bhattacharya, R., and Keith, K. (2024).
Proximal causal inference with text data.
Advances in Neural Information Processing Systems, 37:135983–136017.

Egami, N., Fong, C. J., Grimmer, J., Roberts, M. E., and Stewart, B. M. (2022).
How to make causal inferences using texts.
Science Advances, 8(42):eabg2652.

Fong, C. and Grimmer, J. (2023).
Causal inference with latent treatments.
American Journal of Political Science, 67(2):374–389.

Fu, J. and Green, D. P. (2025).
Causal inference for experiments with latent outcomes: Key results and their implications for design and analysis.
arXiv preprint arXiv:2505.21909.