Marginal effects for time-fixed treatments

Erica E. M. Moodie

Department of Epidemiology, Biostatistics, & Occupational Health
McGill University
Montreal, QC, Canada

erica.moodie@mcgill.ca

Session goals

- When and why is a marginal effect desirable?
- How can the parameters of a marginal effect be estimated?
  - Can we use traditional approaches?
  - What are the 'newer' approaches?
- What assumptions do we need, and how can we check them?

Road map

1. Marginal effects in a point-source treatment setting
   - Definition
   - Regression and stratification
2. The propensity score: a method to recover covariate balance
   - Definition
   - PS stratification
   - PS matching
   - PS regression
   - inverse probability of treatment weighting
3. Assessing confounder balance
4. Double robustness

Concept: Average Potential Outcomes

The *causal* (unconfounded) *effect* of exposure Z on outcome Y is a measure of how much Y changes as Z is manipulated.

- Here Z is not treated as a random variable, but a manipulable quantity that may influence Y.
- Other variables (confounders), X, may also influence Y.
- $Y(z)$ denotes the outcome if the exposure Z is set equal to z:
  - $Y(z)$ is termed a counterfactual or potential outcome.
- A causal quantity of interest is then

$$\mathbb{E}[Y(z)] = \int y f_{Y(z),X}(y,x)\, \mathrm{d}y\mathrm{d}x$$

that is, an average potential outcome (APO).

Aim

Estimate $\mathbb{E}[Y(z)]$ using a random sample of data

$$(x_i, z_i, y_i), i = 1, \ldots, n$$

for z in some set of values

- $z \in \{0, 1\}$
- $z \in \{0, 1, 2, \ldots, K\}$
- $z \in (a, b)$

Small problem: assumptions

Often, we do not have access to experimental data. There is no intervention on behalf of the researcher, the data are recorded observationally.

If we could correctly specify the model $f_{Y|Z,X}(y|z,x)$, or at least the conditional expectation

$$\mathbb{E}[Y|Z=z,X=x]$$

then this would not be a problem, as we could simply use the iterated expectation result and estimate

$$\widehat{\mathbb{E}}[Y(z)] = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[Y|Z=z,X=x_i].$$

Assumptions?

What assumptions do we need to get the 'right' answer, i.e. an unbiased estimator of the marginal mean $\mathbb{E}[Y(z)]$, via regression when data are obtained observationally?

- Correct model specification (of mean of Y given Z and X)
- No unmeasured confounding $\rightarrow$ exchangeability
- Independence $\rightarrow$ no interference
- No extrapolation $\rightarrow$ positivity
- Well-defined exposure $\rightarrow$ cannot have multiple versions of treatment

What if we *cannot* satisfy the first assumption?

Unconfounded effect estimation by design

- We can sometimes estimate the APO (or a contrast of APOs, such as the average treatment effect: ATE) by designing a randomized control trial.
- Recall the setting in the case of a binary exposure:
  - obtain a random sample of size n of individuals from the target population, and measure their X values;
  - according to some random assignment procedure, intervene to assign treatment Z to individuals, and measure their outcome Y;
  - the link between X and Z is broken by the random allocation.

Unconfounded effect estimation by design

- Recall that this procedure led to the valid use of the estimators of the ATE based on (1) and (2) from the previous section.
- The important feature of the randomized study is that we have, for confounders X (indeed all predictors)

$$f_{X|Z}(x|1) \equiv f_{X|Z}(x|0) \quad \text{for all } x,$$

or equivalently, in the case of a binary confounder,

$$\Pr[X = 1|Z = 1] = \Pr[X = 1|Z = 0].$$

- The distribution of X is *balanced* across the two exposure groups; this renders direct comparison of the outcomes possible. Probabilistically, X and Z are independent.

Constructing a balanced sample

- In a non-randomized study, there is a possibility that the two exposure groups are *not balanced*

$$f_{X|Z}(x|1) \neq f_{X|Z}(x|0) \quad \text{for some } x,$$

or in the binary case

$$\Pr[X = 1|Z = 1] \neq \Pr[X = 1|Z = 0].$$

- If X influences Y also, then this imbalance renders direct comparison of outcomes in the two groups impossible.

Constructing a balanced sample

- While *global* balance may not be present, it may be that *local* balance, i.e. within certain strata of the sample, may be present.
- That is, for $x \in \mathcal{S}$ say, we might have balance; within $\mathcal{S}$, X is independent of Z.

$$f_{X|Z:\mathcal{S}}(x|1 : x \in \mathcal{S}) = f_{X|Z}(x|0 : x \in \mathcal{S})$$

- Then, for individuals who have X values in $\mathcal{S}$, there is the possibility of direct comparison of the treated and untreated groups.
- We might then restrict attention to causal statements relating to stratum $\mathcal{S}$.

Constructing a balanced sample

- For discrete confounders, we might consider defining strata where the X values are *precisely matched*, and then comparing treated and untreated within those strata.
- Consider matching strata $\mathcal{S}_1, \ldots, \mathcal{S}_K$. We would then be able to compute the ATE by noting that

$$\mathbb{E}[Y(1) - Y(0)] = \sum_{k=1}^{K} \mathbb{E}[Y(1) - Y(0) | X \in \mathcal{S}_k] \Pr[X \in \mathcal{S}_k]$$

  - $\mathbb{E}[Y(1) - Y(0) | X \in \mathcal{S}_k]$ may be estimated nonparametrically from the data by using (1) or (2) for data restricted to have $x \in \mathcal{S}_k$.
  - $\Pr[X \in \mathcal{S}_k]$ may be estimated using the empirical proportion of x that lie in $\mathcal{S}_k$.

Constructing a balanced sample

- For continuous confounders, we might consider the same strategy: consider matching strata $\mathcal{S}_1, \ldots, \mathcal{S}_K$. Then the formula

$$\mathbb{E}[Y(1) - Y(0)] = \sum_{k=1}^{K} \mathbb{E}[Y(1) - Y(0) | X \in \mathcal{S}_k] \Pr[X \in \mathcal{S}_k]$$

  still holds.
- However we must assume a model for how $\mathbb{E}[Y(1) - Y(0) | X \in \mathcal{S}_k]$ varies with x for $x \in \mathcal{S}_k$.
- In both cases, inference is restricted to the set of X space contained in

$$\bigcup_{k=1}^{K} \mathcal{S}_k.$$

Constructing a balanced sample

- In the continuous case, the above calculations depend on the assumption that the treatment effect is similar for x values that lie 'close together' in predictor (confounder) space. However
  I. Unless we can achieve exact matching, then the term 'close together' needs careful consideration.
  II. If X is moderate or high-dimensional, there may be insufficient data to achieve adequate matching to facilitate the estimation of the terms
  $$\mathbb{E}[Y(1) - Y(0)|X \in \mathcal{S}_k];$$
  recall that we need a large enough sample of treated and untreated subjects in stratum $\mathcal{S}_k$.
- Nevertheless, matching in this fashion is an important tool in causal comparison.

Balance via the propensity score

- We now come to the important concept of the propensity score that facilitates causal comparison via a balancing approach.
- Recall: our goal is to mimic the construction of the randomized study that facilitates direct comparison between treated and untreated groups.
- We may not be able to achieve this globally, but possibly can achieve it locally in strata of X space.

- The question is how to define these strata.

Balance via the propensity score

- Recall that in the binary exposure case, balance corresponds to being able to state that within $\mathcal{S}$, X is *independent* of Z.
- This can be achieved if $\mathcal{S}$ is defined in terms of a statistic, $e(X)$ say. That is, we consider the conditional distribution

$$f_{X|Z,e(X)}(x|z,e)$$

and attempt to ensure that, given $e(X) = e$, Z is independent of X, so that within strata of $e(X)$, the treated and untreated groups are directly comparable.

Balance via the propensity score

- By Bayes theorem, for $z = 0, 1$, we have that

$$f_{X|Z,e(X)}(x|z,e) = \frac{f_{Z|X,e(X)}(z|x,e) f_{X|e(X)}(x|e)}{f_{Z|e(X)}(z|e)} \tag{1}$$

- Now, as Z is binary, we must be able to write the density in the denominator as

$$f_{Z|e(X)}(z|e) = p(e)^z (1 - p(e))^{1-z} \qquad z \in 0, 1$$

where $p(e)$ is a probability, a function of the fixed value e, and where $0 < p(e) < 1$.

Balance via the propensity score

- Therefore, in order to make the density $f_{X|Z,e(X)}(x|z,e)$ functionally independent of z, and so achieve the independence we seek, we need

$$f_{Z|X,e(X)}(z|x,e) = p(e)^z(1-p(e))^{1-z} \qquad z \in 0,1.$$

- But $e(X)$ is a function of X, so automatically we have that

$$f_{Z|X,e(X)}(z|x,e) \equiv f_{Z|X}(z|x).$$

Therefore, we require that

$$f_{Z|X}(z|x) = f_{Z|X}(z|x,e) = p(e)^z(1-p(e))^{1-z} \equiv f_{Z|e(X)}(z|e)$$

for all relevant z, x.

Balance via the propensity score

- This can be achieved by choosing the statistic

$$e(x) = \Pr_{Z|X}[Z = 1|x]$$

and setting $p(.)$ to be the identity function, so that

$$f_{Z|X}(z|x) = e^z(1-e)^{1-z} \quad z = 0, 1.$$

- More generally, choosing $e(x)$ to be some monotone transform of $\Pr_{Z|X}[Z = 1|x]$ would also achieve the same balance.
- The corresponding random variable $e(X)$ defines the strata via which the causal calculation can be considered.

Balance via the propensity score

- The function $e(x)$ defined in this way is the *propensity score*[1]. It has the following important properties
  - (i) as seen above, it is a balancing score; conditional on $e(X)$, X and Z are independent.
  - (ii) it is a *scalar* quantity, irrespective of the dimension of X.
  - (iii) in noting that for balance we require that

$$f_{Z|X}(z|x) \equiv f_{Z|e(X)}(z|e),$$

the above construction demonstrates that if $\widetilde{e}(X)$ is another balancing score, then $e(X)$ is a function of $\widetilde{e}(X)$. That is, $e(X)$ is the 'coarsest' balancing score.

[1] see Rosenbaum & Rubin (1983), Biometrika

Evaluating the propensity score

- To achieve balance we must have

$$e(X) = \Pr[Z = 1|X]$$

correctly specified; that is, for confounders X, we must precisely specify the model $\Pr[Z = 1|X]$.

  - If X comprises entirely discrete components, then we may be able to estimate $\Pr[Z = 1|X]$ entirely nonparametrically, and satisfactorily if the sample size is large enough.
  - If X has continuous components, it is common to use parametric modelling, with

$$e(X; \alpha) = \Pr[Z = 1|X; \alpha].$$

Balance then depends on *correct specification* of this model.

Unconfoundedness given the propensity score

- The assumption of 'no unmeasured confounders' amounts to assuming that the potential outcomes are jointly *independent* of exposure assignment given the confounders, that is

$$\{Y(0), Y(1)\} \perp\!\!\!\perp Z \mid X.$$

- With a correctly specified propensity score, we now have that

$$Y(z) \perp\!\!\!\perp Z \mid e(X) \qquad \text{for all } z.$$

Estimation using the propensity score

- We now consider the same stratified estimation strategy as before, but using $e(X)$ instead X to stratify.
- Consider strata $\mathcal{S}_1, \ldots, \mathcal{S}_K$ defined via $e(X)$. In this case, recall that

$$0 < e(X) < 1$$

  so we might consider an equal quantile partition, say using quintiles.
- Then we have

$$\mathbb{E}[Y(1) - Y(0)] = \sum_{k=1}^{K} \mathbb{E}[Y(1) - Y(0) | e(X) \in \mathcal{S}_k] \Pr[e(X) \in \mathcal{S}_k]$$

  still holds approximately if the $\mathcal{S}_k$ are small enough.

Estimation using the propensity score

- This still requires us to be able to estimate

$$\mathbb{E}[Y(1) - Y(0)|e(X) \in \mathcal{S}_k]$$

  which requires us to have a sufficient number of treated and untreated individuals with $e(X) \in \mathcal{S}_k$ to facilitate the 'direct comparison' within this stratum.
- If the expected responses are constant across the stratum, the formulae (1) and (2) may be used.

Matching

The derivation of the propensity score indicates that it may be used to construct *matched* individuals or groups that can be compared directly.

- if two individuals have precisely the same value of $e(x)$, then they are exactly matched;
- if one of the pair is treated and the other untreated, then their outcomes can be compared directly, as any imbalance between their measured confounder values has been removed by the fact that they are matched on $e(x)$;
- this is conceptually identical to the standard procedure of matching in two-group comparison.

Matching

For an exactly matched pair (i_1, i_0), treated and untreated respectively, the quantity

$$Y_{i_1} - Y_{i_0}$$

is an unbiased estimate of the ATE

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

more typically we might choose m such matched pairs, usually with different $e(x)$ values across pairs, and use the estimate

$$\frac{1}{m} \sum_{i=1}^{m} (Y_{i_1} - Y_{i_0})$$

Matching

Exact matching is difficult to achieve, therefore we more commonly attempt to achieve approximate matching:

- may match one treated to M untreated ($1 : M$ matching)
- caliper matching;
- nearest neighbour/kernel matching;
- matching with replacement.

Most standard software packages have functions that provide automatic matching using a variety of methods.

Propensity Score Regression

Up to this point we have considered using the propensity score for stratification, that is, to produce directly comparable groups of treated and untreated individuals.

Causal comparison can also be carried out using regression techniques: that is, we consider building an estimator of the APO by *regressing* the outcome on a function of the exposure and the propensity score.

Regressing on the propensity score is a means of controlling the confounding.

Propensity Score Regression

If we construct a model

$$\mathbb{E}[Y|Z = z, e(Z, X) = e] = \mu(z, e)$$

then because potential outcomes $Y(z)$ and Z are independent given $e(Z, X)$, we have

$$\mathbb{E}[Y(z)|e(Z, X) = e] = \mathbb{E}[Y|Z = z, e(z, X) = e] = \mu(z, e)$$

and therefore

$$\mathbb{E}[Y(z)] = \mathbb{E}_{e(z,X)}[\mathbb{E}[Y|Z = z, e(z, X)]] = \mathbb{E}_{e(z,X)}[\mu(z, e(z, X))].$$

Propensity Score Regression

That is, to estimate the APO, we might

- fit the propensity score model $e(Z, X)$ to the observed exposure and confounder data by regressing Z on X;
- fit the conditional outcome model $\mu(z, e)$ using the fitted $e(Z, X)$ values, $\widehat{e}(z_i, x_i)$;
- for each z of interest, estimate the APO by

$$\frac{1}{n}\sum_{i=1}^{n}\widehat{\mu}(z, \widehat{e}(z, x_i)).$$

Propensity Score Regression

If the propensity function $e(Z, X) \equiv e(X)$, we proceed similarly, and construct a model

$$\mathbb{E}[Y|Z = z, e(X) = e] = \mu(z, e)$$

then

$$\mathbb{E}[Y(z)|e(X) = e] = \mathbb{E}[Y|Z = z, e(X) = e] = \mu(z, e)$$

and therefore

$$\mathbb{E}[Y(z)] = \mathbb{E}_{e(X)}[\mathbb{E}[Y|Z = z, e(X)]] = \mathbb{E}_{e(X)}[\mu(z, e(X))].$$

Propensity Score Regression

To estimate the APO:

- fit the propensity score model $e(X)$ to the observed exposure and confounder data by regressing Z on X;
- fit the conditional outcome model $\mu(z, e)$ using the fitted $e(X)$ values, $\widehat{e}(x_i)$;
- for each z of interest, estimate the APO by

$$\frac{1}{n}\sum_{i=1}^{n}\widehat{\mu}(z, \widehat{e}(x_i)).$$

Example: Binary Exposure

We specify

- $e(X;\alpha) = \Pr[Z = 1|X, \alpha]$ then regress Z on X to obtain $\widehat{\alpha}$ and fitted values $\widehat{e}(X) \equiv e(X;\widehat{\alpha})$.
- $\mathbb{E}[Y|Z = z, e(X) = e;\beta] = \mu(z, e;\beta)$ and estimate this model by regressing y_i on z_i and $\widehat{e}(x_i)$. For example, we might have that

$$\mathbb{E}[Y|Z = z_i, e(X_i) = e_i;\beta] = \beta_0 + \beta_1 z_i + \beta_2 e_i.$$

  This returns $\widehat{\beta}$.

We finally compute the predictions under this model, and average them to obtain the APO estimate

$$\widehat{\mathbb{E}}[Y(z)] = \frac{1}{n}\sum_{i=1}^{n} \mu(z, \widehat{e}(x_i);\widehat{\beta}).$$

Example: Continuous Exposure

In the case of a continuous exposure, we have a parametric probability density for the exposure

$$e(Z, X; \alpha) = f_{Z|X}(Z|X; \alpha)$$

for which we estimate α by regressing Z on X to obtain $\widehat{\alpha}$ and fitted values $\widehat{e}(Z, X) \equiv e(Z, X; \widehat{\alpha})$.

Then we specify outcome model

$$\mathbb{E}[Y|Z = z, e(X) = e; \beta] = \mu(z, e; \beta)$$

and estimate this model by regressing y_i on z_i and $\widehat{e}(z_i, x_i)$. Again, we might have that

$$\mathbb{E}[Y|Z = z_i, e(Z_i, X_i) = e_i; \beta] = \beta_0 + \beta_1 z_i + \beta_2 e_i.$$

This returns $\widehat{\beta}$.

Example: Binary Exposure

We then compute the predictions under this model, and average them to obtain the APO estimate

$$\widehat{\mathbb{E}}[Y(z)] = \frac{1}{n}\sum_{i=1}^{n} \mu(z, \widehat{e}(z, x_i); \widehat{\beta}).$$

Note that here the propensity terms that enter into μ are computed at the target z values, and

not the observed exposure values.

Propensity Score Regression

These procedures require us to make two modelling choices:

- the propensity model, $e(Z, X)$ or $e(X)$;
- the outcome mean model $\mu(z, b)$.

Unfortunately, *both models must be correctly specified* for consistent inference.

Misspecification of the outcome mean model will lead to bias; this model needs to capture the outcome to exposure and propensity function relationship correctly.

Average potential outcome

If we could intervene at the population level to set $Z = z$ for all individuals independently of their X value, we might rewrite $\mathbb{E}[Y(z)]$ as

$$\mathbb{E}[Y(z)] = \int y \mathbb{1}_z(z) f_{Y(z),X}(y, x) \, \mathrm{d}y \, \mathrm{d}z \, \mathrm{d}x$$

and take a random sample from the population with density

$$\mathbb{1}_z(z) f_{Y(z),X}(y, x) \equiv \mathbb{1}_z(z) f_{Y|Z,X}(y|z, x) f_X(x).$$

We could then construct the moment estimate

$$\widehat{\mathbb{E}}[Y(z)] = \frac{1}{n} \sum_{i=1}^{n} y_i$$

as $z_i = z$ for all i.

Average potential outcome: Experimental data

In a randomized (experimental) study, suppose that exposure $Z = z$ is assigned with probability determined by $f_Z(z)$. Then

$$\mathbb{E}[Y(z)] = \frac{\int y \, \mathbb{1}_z(z) \, f_{Y|Z,X}(y|z,x) f_X(x) f_Z(z) \, \mathrm{d}y \, \mathrm{d}z \, \mathrm{d}x}{\int \mathbb{1}_z(z) f_Z(z) \, \mathrm{d}z}$$

This suggests the Monte Carlo estimates

$$\widehat{\mathbb{E}}[Y(z)] = \frac{\sum_{i=1}^{n} \mathbb{1}_z(z_i) y_i}{\sum_{i=1}^{n} \mathbb{1}_z(z_i)} \qquad \text{or} \qquad \widehat{\mathbb{E}}[Y(z)] = \frac{1}{n f_Z(z)} \sum_{i=1}^{n} \mathbb{1}_z(z_i) y_i$$

Average potential outcome

Commonly, we want to carry out a comparison of average potential outcomes at different values of z, e.g.:

$$\widehat{\mathbb{E}}[Y(1) - Y(0)] = \frac{\sum_{i=1}^{n} \mathbb{1}_{z=1}(z_i) y_i}{\sum_{i=1}^{n} \mathbb{1}_{z=1}(z_i)} - \frac{\sum_{i=1}^{n} \mathbb{1}_{z=0}(z_1) y_i}{\sum_{i=1}^{n} \mathbb{1}_{z=0}(z_i)}$$

or

$$\widehat{\mathbb{E}}[Y(1) - Y(0)] = \frac{1}{n f_Z(z=1)} \sum_{i=1}^{n} \mathbb{1}_{z=1}(z_i) y_i - \frac{1}{n f_Z(z=0)} \sum_{i=1}^{n} \mathbb{1}_{z=0}(z_i) y_i.$$

Average potential outcome: Observational data

Denote by $P_{\mathcal{E}}$ the probability measure for samples drawn under the experimental measure corresponding to the density

$$f^{\mathcal{E}}_{Y|Z,X}(y|z,x)f^{\mathcal{E}}_{X}(x)f^{\mathcal{E}}_{Z}(z)$$

Now consider the case where the data arise from the observational (non-experimental) measure $P_{\mathcal{O}}(\mathrm{d}y, \mathrm{d}z, \mathrm{d}x)$. We have

$$\mathbb{E}[Y(z)] = \frac{1}{f^{\mathcal{E}}_{Z}(z)} \int y \mathbb{1}_{z}(z)\, P_{\mathcal{E}}(\mathrm{d}y, \mathrm{d}z, \mathrm{d}x)$$

$$= \frac{1}{f^{\mathcal{E}}_{Z}(z)} \int y \mathbb{1}_{z}(z)\, \underbrace{\frac{P_{\mathcal{E}}(\mathrm{d}y, \mathrm{d}z, \mathrm{d}x)}{P_{\mathcal{O}}(\mathrm{d}y, \mathrm{d}z, \mathrm{d}x)}}_{①} P_{\mathcal{O}}(\mathrm{d}y, \mathrm{d}z, \mathrm{d}x)$$

In terms of densities ① becomes

$$\frac{f^{\mathcal{E}}_{Y|Z,X}(y|z,x)f^{\mathcal{E}}_{Z}(z)f^{\mathcal{E}}_{X}(x)}{f^{\mathcal{O}}_{Y|Z,X}(y|z,x)f^{\mathcal{O}}_{Z|X}(z|x)f^{\mathcal{O}}_{X}(x)} = \frac{f^{\mathcal{E}}_{Y|Z,X}(y|z,x)}{f^{\mathcal{O}}_{Y|Z,X}(y|z,x)} \times \frac{f^{\mathcal{E}}_{Z}(z)}{f^{\mathcal{O}}_{Z|X}(z|x)} \times \frac{f^{\mathcal{E}}_{X}(x)}{f^{\mathcal{O}}_{X}(x)}$$

Estimation

This suggests the (nonparametric) estimators

$$\widehat{\mathbb{E}}[Y(z)] = \frac{1}{n}\sum_{i=1}^{n} \frac{\mathbb{1}_z(Z_i)Y_i}{f^{\mathcal{O}}_{Z|X}(Z_i|X_i)} \qquad \text{(IPW0)}$$

which is unbiased, or

$$\widehat{\mathbb{E}}[Y(z)] = \frac{\sum_{i=1}^{n} \frac{\mathbb{1}_z(Z_i)Y_i}{f^{\mathcal{O}}_{Z|X}(Z_i|X_i)}}{\sum_{i=1}^{n} \frac{\mathbb{1}_z(Z_i)}{f^{\mathcal{O}}_{Z|X}(Z_i|X_i)}} \qquad \text{(IPW)}$$

which is consistent, each provided $f^{\mathcal{O}}_{Z|X}(.|.)$ correctly specifies the conditional density of Z given X for all (z, x).

The inverse probability weighting constructs a pseudo-population in which there are no imbalances on measured covariates between the exposure groups.

Assumptions

What assumptions do we need to get the 'right' answer, i.e. an unbiased estimator of the marginal mean $\mathbb{E}[Y(z)]$, via IPW?

- Correct model specification (of mean of Z given X)
- No unmeasured confounding
- Independence
- No extrapolation
- Well-defined exposure

Parametric modelling: two-stage approach

In the formulation, the nonparametric models

$$f^{\mathcal{O}}_{Z|X}(z|x) \qquad \mu(z,x)$$

are commonly replaced by parametric models

$$f^{\mathcal{O}}_{Z|X}(z|x;\alpha) \qquad \mu(z,x;\beta) = \int y f^{\mathcal{O}}_{Y|Z,X}(y|z,x;\beta)\,\mathrm{d}y.$$

Parameters (α, β) are estimated from the observed data by regressing

- Stage I: Z on X using $(z_i, x_i), i = 1, \ldots, n$.
- Stage II: Y on (Z, X) using $(y_i, z_i, x_i), i = 1, \ldots, n$.

and using plug-in version of (IPW).

Confounder balance

- In PS-based methods, the goal of the treatment model is to eliminate imbalance in the distribution of covariates between treatment and untreated subjects.

- Some measures of balance:
  - Standardized mean difference or proportion:

$$\frac{\bar{x}^{1,sw} - \bar{x}^{0,sw}}{\sqrt{0.5(v^{1,sw} + v^{0,sw})}}$$

    where $\bar{x}^{z,w} = \frac{1}{n}\sum_{i=1}^{n} \frac{\mathbb{1}_z(Z_i)X_i}{f^{\mathcal{O}}_{Z|X}(Z_i|X_i)}$, i.e. the weighted sample mean of variable X among those with treatment value z (for binary treatment), and similarly $v^{z,w}$ is the weighted variance estimate.
  - Visual examination of weighted empirical CDFs among the treated and untreated (for binary or categorical treatment).

Confounder balance

- Some methods to avoid:
  - C-statistic.
  - Significance tests.

- This ties in with the key point that the goal is to eliminate imbalance and thereby remove the effects of confounding.
- The goal is *not* to build an excellent predictive model for the treatment.

Assessing balance: example

- In this example, we will explore propensity score based analyses using the publicly available (U.S.) National Health and Nutrition Examination Survey (NHANES). For this, I installed `NHANES`, `tableone`, and `Matching` in `R`.

- We will focus our analysis on the question of whether currently smoking affects average systolic blood pressure. The variables we will need are: `BPSysAve`, `SmokeNow`, `Gender`, `Age`, `Race3`, `Education`, `MaritalStatus`, and `Poverty` where the first two are the outcome and exposure of interest and the remaining are potential confounders.
- Additionally, we will restrict our attention to adults ($>$ 17 years old) in the second wave of the survey.

Assessing balance: example

```
> library(NHANES)
> library(tableone)
> library(Matching)
>
> NHANES$SmokeNow <- as.numeric(NHANES$SmokeNow)-1
> small.nhanes <- na.omit(NHANES[NHANES$SurveyYr=="2011_12"
    & NHANES$Age > 17,c(3,4,8:11,13,25,61)])
> dim(small.nhanes) ## 1377
>
> vars <- c("Gender", "Age", "Race3", "Education",
    "MaritalStatus", "Poverty")
> tabUnmatched <- CreateTableOne(vars = vars,
    strata = "SmokeNow", data = small.nhanes,
    test = FALSE)
```

Assessing balance: example

```
> print(tabUnmatched, smd = TRUE)
                     Stratified by SmokeNow
                      0              1              SMD
  n                     782            595
  Gender = male (%)     432 (55.2)     369 (62.0)    0.138
  Age (mean (sd))     54.33 (16.52)  44.96 (15.11)   0.592
  Race3 (%)                                         0.315
     Asian              25 ( 3.2)      15 ( 2.5)
     Black              43 ( 5.5)      64 (10.8)
     Hispanic           26 ( 3.3)      38 ( 6.4)
     Mexican            45 ( 5.8)      35 ( 5.9)
     White             630 (80.6)     416 (69.9)
     Other              13 ( 1.7)      27 ( 4.5)
```

Assessing balance: example

Education (%)			0.512
8th Grade	59 (7.5)	33 (5.5)	
9 - 11th Grade	71 (9.1)	120 (20.2)	
High School	152 (19.4)	151 (25.4)	
Some College	256 (32.7)	210 (35.3)	
College Grad	244 (31.2)	81 (13.6)	
MaritalStatus (%)			0.488
Divorced	85 (10.9)	77 (12.9)	
LivePartner	61 (7.8)	96 (16.1)	
Married	453 (57.9)	240 (40.3)	
NeverMarried	108 (13.8)	142 (23.9)	
Separated	6 (0.8)	14 (2.4)	
Widowed	69 (8.8)	26 (4.4)	
Poverty (mean (sd))	3.11 (1.65)	2.38 (1.58)	0.453

Assessing balance: example

Assessing balance – original sample eCDFs in smokers and non-smokers for age:

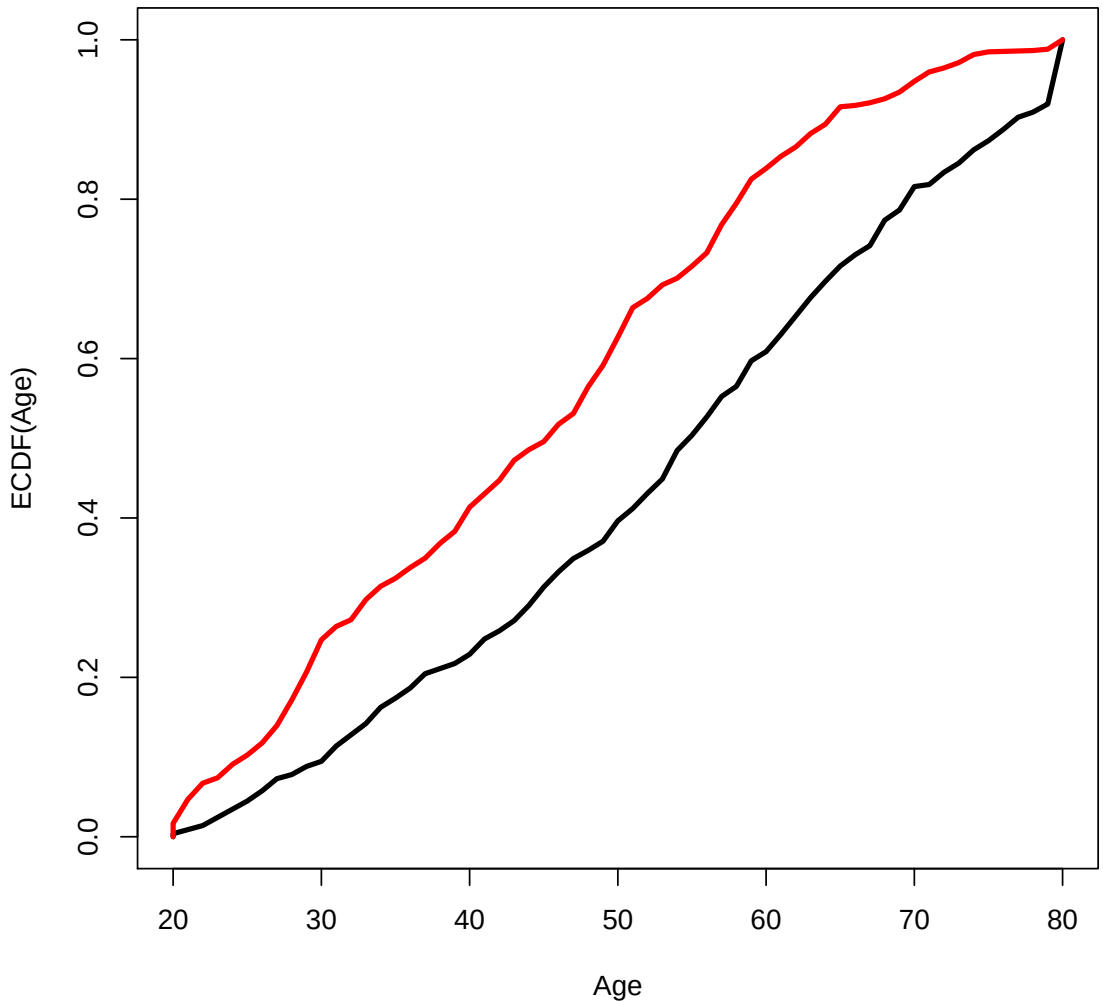

Assessing balance: example

Table. Standardized mean differences: NHANES.

Var.	PS Quintiles						
	Orig.	Q1	Q2	Q3	Q4	Q5	Match	IPW
Gender	0.138	0.102	0.104	0.029	0.200	0.031	0.006	0.023
Age	0.592	0.257	0.171	0.099	0.311	0.164	0.002	0.014
Race	0.315	0.317	0.112	0.344	0.415	0.287	0.120	0.052
Educ.	0.512	0.538	0.417	0.280	0.238	0.302	0.133	0.029
Marital	0.488	0.432	0.239	0.272	0.233	0.261	0.094	0.023
Poverty	0.453	0.087	0.126	0.114	0.004	0.146	0.049	0.000

Assessing balance: example

Assessing balance – eCDFs within quintiles of PS in smokers and non-smokers for age:

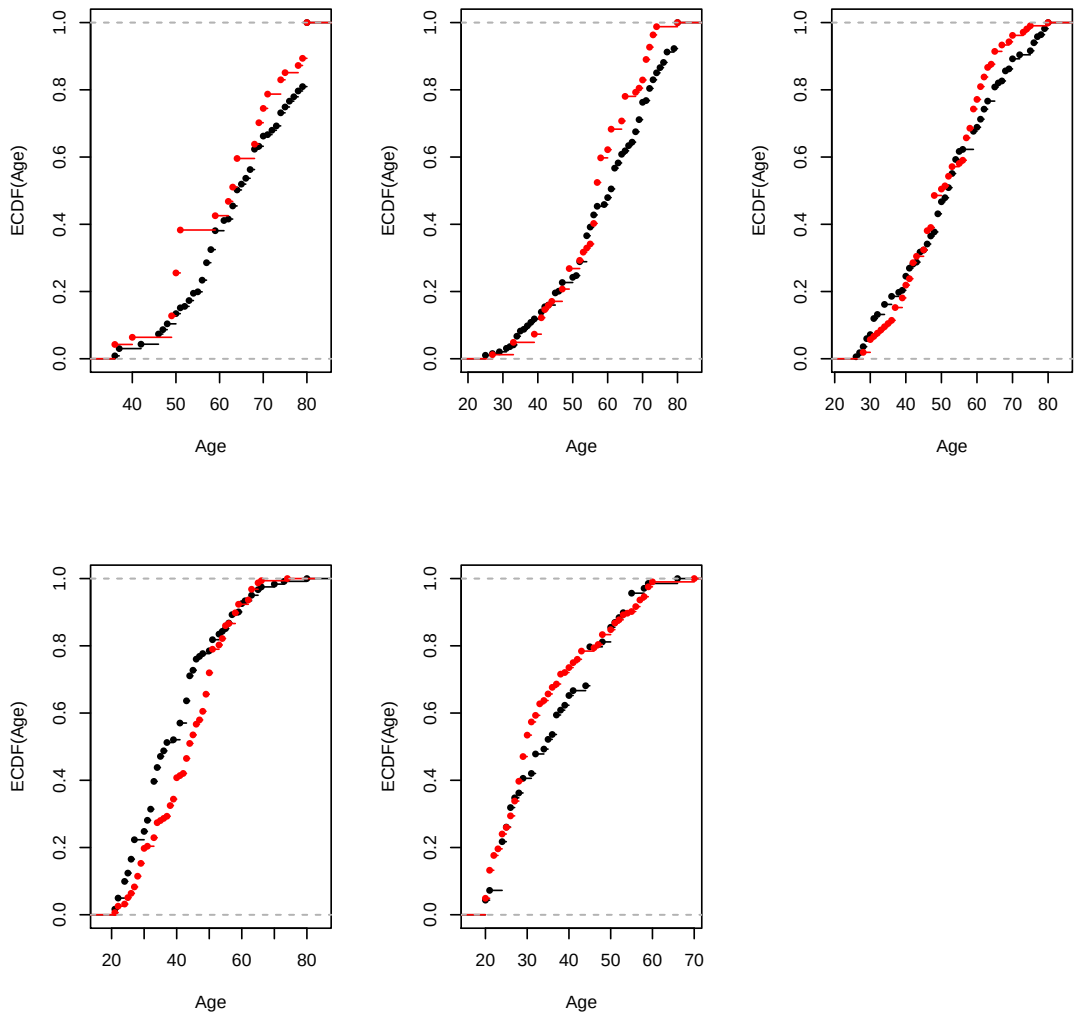

Assessing balance: example

Assessing balance – eCDFs in smokers and non-smokers for age, matched and IPW:

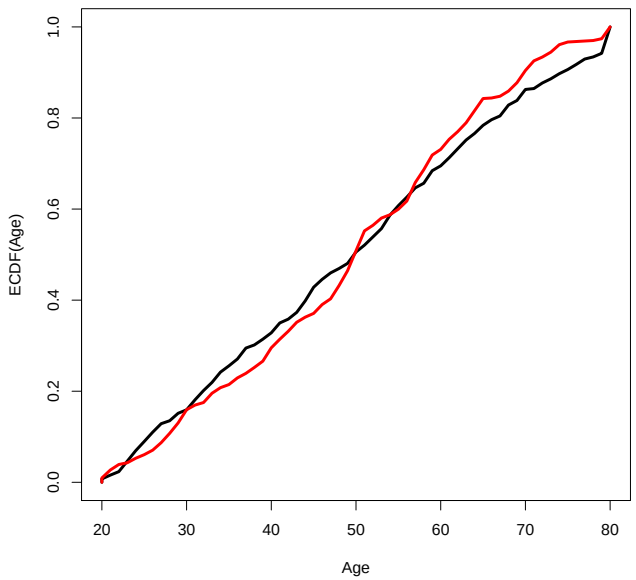

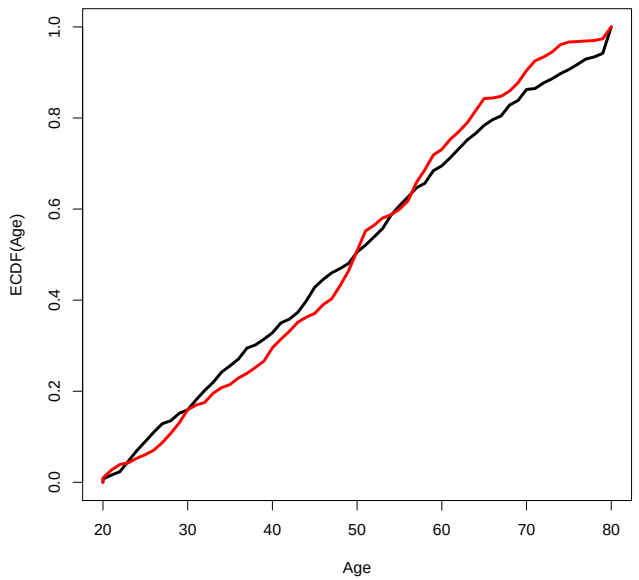

Double robustness

- The IPW is popular, perhaps unduly so given that it is provably less efficient than PS regression.
- Can we improve upon it?

Double robustness

The IPW can be *augmented*. Note that

$$\mathbb{E}[Y(z)] = \mathbb{E}[Y(z) - \mu(z, X)] + \mathbb{E}[\mu(z, X)]$$

where $\mu(z, x) = \mathbb{E}[Y|Z = z, X = x]$.

This gives the alternate estimator

$$\widehat{\mathbb{E}}[Y(z)] = \frac{1}{n}\sum_{i=1}^{n} \frac{\mathbb{1}_z(Z_i)(Y_i - \mu(Z_i, X_i))}{f^{\mathcal{O}}_{Z|X}(Z_i|X_i)} + \frac{1}{n}\sum_{i=1}^{n} \mu(z, X_i).$$

(AIPW)

Doubly robust IPW

Equation (AIPW) is doubly robust (i.e. consistent even if one of $f^{\mathcal{O}}_{Z|X}(z|x)$ and $\mu(z,x)$ is misspecified).

- If $\mu(z, X_i)$ is correctly specified, then $Y_i - \mu(Z_i, X_i) \to 0$, and the first term in the augmented estimator disappears (asymptotically), leaving the term $\frac{1}{n}\sum_{i=1}^n \mu(z, X_i)$ which is consistent for $\mathbb{E}[Y(z)]$.
- If $f^{\mathcal{O}}_{Z|X}(Z_i|X_i)$ is correctly specified, then $\frac{\mathbb{1}_z(Z_i)}{f^{\mathcal{O}}_{Z|X}(Z_i|X_i)} \to 1$, and so

$$\frac{1}{n}\sum_{i=1}^n \frac{\mathbb{1}_z(Z_i)(-\mu(Z_i,X_i))}{f^{\mathcal{O}}_{Z|X}(Z_i|X_i)} + \frac{1}{n}\sum_{i=1}^n \mu(z, X_i) \to 0,$$

  leaving $\frac{1}{n}\sum_{i=1}^n \frac{\mathbb{1}_z(Z_i)Y_i}{f^{\mathcal{O}}_{Z|X}(Z_i|X_i)}$.

Further, $\text{Var}_{\text{AIPW}} \leq \text{Var}_{\text{IPW}}$.

Time-fixed treatments & causal inference: Summary

- Causal estimation of time-fixed treatment effects can be achieved through traditional approaches such as regression or stratification.
- The propensity score is simply a model for the exposure given confounding variables; note that 'instruments' (variables that only predict treatment) should be omitted.
- The PS can be used in many ways: stratification, matching, regression, IPW. Regression is most efficient, but it is prudent to be flexible.
- Balance can be assessed visually and through tables.

Key points: Summary

- A marginal summary attempts to answer questions relevant to policy makers: *what is the expected outcome, averaged over the covariate distribution in my population?*
- Such questions help to avoid the 'trap' of contrasting those who are observed to be treated and untreated, as these may be very different (w.r.t confounding variables) groups of individuals.
- To recover a marginal summary, we need to restore, or create, balance on covariates between the treatment groups.
- We can only restore balance on covariates that we have measured. It is crucial to understand the context of the question to begin to assess whether all confounders have been measured.

Acknowledgements

Collaborator: David Stephens (McGill)

Selected references:

- Rosenbaum and Rubin (1983) The central role of the propensity score in observational studies for causal effects. *Biometrika* **6**:41–55.
- Rubin (1978) Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics* **6**:34–58.