

Soc. Stats Reading Group Four: Sensitivity Analysis


Rebecca Johnson

October 27th, 2016

Purpose of sensitivity analysis, to deal with conversations like these...

- ▶ You: Here are the results of my matching (*or insert other method for causal inference with observational data that relies on some form of ignorability assumption*) analysis showing that becoming unemployed causes a higher likelihood of opioid use
- ▶ Your skeptical adviser: when estimating your propensity score, did you include a measure of whether the person worked in a manual occupation on the losing end of skill-based technological change?
- ▶ You: Of course!
- ▶ Your skeptical adviser: what about a measure of a person's (pre-treatment) degree of existential suffering?
- ▶ You: I think that's impossible to observe or measure...
- ▶ Your skeptical adviser: well how large would a difference in inherent potential for opioid use among those more likely to become unemployed need to be to bias your results? When would it lead you to find that becoming unemployed causes a *lower* likelihood of opioid use?
- ▶ You: I have no idea...I guess I need to learn about sensitivity analysis...

In diagram form



Outline

► Blackwell

1. Review of causal quantities of interest in potential outcomes framework
2. Brief diversion to Rosenbaum (2002) to build intuition about purpose behind sensitivity analysis
3. Blackwell's approach: "de-confound" the dependent variable using a confounding function
4. Illustrate approach with (relatively) basic case
5. Focus on three extensions of basic case:
 - 5.1 Change confounding function from one-sided bias to alignment bias
 - 5.2 Re-parametrize confounding function to express magnitude of confounding in terms of variation in R^2 rather than difference in mean inherent potential outcomes between treatment and control (α)
 - 5.3 Going from static, cross-sectional treatment assignment to treatments over time (dynamic case)

- ## ► Throughout, briefly contextualize with other approaches discussed in Morgan and Winship

Causal quantities of interest and underpinning assumptions

Quantities:

- ▶ Average treatment effect (ATE):

$$\tau = E[Y_i(1) - Y_i(0)] = E[Y_i(1)] - E[Y_i(0)]$$

- ▶ Average treatment effect among treated units (ATT):

$$\tau_{att} = E[Y_i(1) - Y_i(0)|A_i = 1] = E[Y_i(1)|A_i = 1] - E[Y_i(0)|A_i = 1]$$

Main assumptions:

- ▶ Consistency (less of a focus here; violated, for instance, by spillover of treatment onto untreated units)
- ▶ Ignorability- version depends on whether we're concerned with estimating ATE versus ATT:
 1. ATE requires ignorability among both treatment and control units, which means for covariates X_i and treatments $a \in (0, 1)$:

$$Y_i(a) \perp A_i | X_i = (Y_i(1), Y_i(0)) \perp A_i | X_i$$

2. ATT requires ignorability only among control:

$$Y_i(0) \perp A_i | X_i$$

Your adviser again...

$$(Y_i(1), Y_i(0)) \perp A_i | X_i$$

1. To satisfy above assumption, can keep on adding X_i to condition on (while paying attention to post-treatment issues discussed in week 1), but in observational data, there will always remain unobserved confounders correlated with both treatment status and potential outcomes (e.g., your adviser's comment about existential suffering)
2. Sensitivity analysis: more systematically explore how the correlation between a unit's probability of receiving treatment and that unit's potential outcomes affects magnitude and direction of estimated treatment effect
3. What's up next:
 - 3.1 More in-depth review of Rosenbaum (2002) than in Blackwell.
Why? Still common form of sensitivity analysis, and also gestures at approach of modeling selection into treatment
 - 3.2 Blackwell article

More background on Rosenbaum (2002)

Illustrating with unemployment and opioid example:¹

1. $\pi_i = Pr(A_i = 1)$; $1 - \pi_i = Pr(A_i = 0)$; X_i are observed covariates
2. Imagine *Bob* and *Jim*, where $X_{Bob} = X_{Jim}$ (e.g., same manufacturing job, same age, same observed disability status). Bob and Jim's odds of treatment (becoming unemployed) are:

$$Odds_{Bob} = \frac{\pi_{Bob}}{1 - \pi_{Bob}}$$

$$Odds_{Jim} = \frac{\pi_{Jim}}{1 - \pi_{Jim}}$$

3. The sensitivity parameter, Γ , is the odds ratio of these two probabilities of treatment, or the odds of Bob being unemployed over the odds of identical observed covariate Jim being unemployed:

$$\Gamma = \frac{\frac{\pi_{Bob}}{1 - \pi_{Bob}}}{\frac{\pi_{Jim}}{1 - \pi_{Jim}}} = \frac{Odds_{Bob}}{Odds_{Jim}}$$

4. While for Blackwell $q(a, x) = 0$ is case where we assume ignorability assumption is satisfied, for Rosenbaum, the case where the *true* value of $\Gamma = 1$ is case where ignorability is satisfied (note that the *observed* value of Γ will always be 1 for two obs. with same observed covariates since all we have are these observed covars to calculate the OR)

¹ Credit to Bertolli (2013) for helpful slides on Rosenbaum bounds

More background on Rosenbaum continued

Basic procedure - engage in a thought experiment where we see how changing Γ to reflect different magnitudes of confounding by unobserved variables affects results:

1. Choose range of Γ that represent worst case scenarios of different odds of Bob versus Jim becoming treated (unemployed) despite same observed covariates (e.g., if you think odds might only differ slightly, so π_{Bob} and π_{Jim} , though not equal, are within the range of 0.33 to 0.66, can calculate Γ range as follows):
 - 1.1 Lower bound: $\frac{0.33}{1-0.33} = 0.5$
 - 1.2 Upper bound: $\frac{0.66}{1-0.66} = 2$
 - 1.3 $0.5 \leq \Gamma \leq 2$ (no individual is more than twice as likely than someone with same covariates to become unemployed)
2. Increment through different values of Γ in that range to see how significance and size of treatment effect changes; estimates for how these quantities change is based on exact test that corresponds to nature of dependent variable
 - 2.1 Binary dependent variable: McNemar's exact test
 - 2.2 Continuous dependent variable: Wilcoxon signed rank test for p value and Hodges-Lehmann for point estimate

Illustrating with binary treatment example

Outcomes for matched pairs:

		<i>Employed</i>	
		Opioid	No opioid
<i>Unemployed</i>	Opioid	10	b = 50
	No opioid	c = 20	140

- ▶ If $\Gamma = 1$ (so assume individuals with same observed covariates have same probability of employed and unemployed), then McNemar's exact p-value is, where $n = b + c$

$$2 * \sum_{i=b}^n \binom{n}{i} 0.5^i 0.5^{n-i} \approx 0.00044$$

- ▶ When you increase Γ , the probabilities in red above are no longer 0.5 (become: $\pi = \frac{\Gamma}{1+\Gamma}$ and $1 - \pi$) and p-value increases
- ▶ Putting into code (mcn.exact.p finds the non-summation part of above expression):

```
mcn.exact.p <- function(n, b, prob){  
  choose(n, b) * prob*b * (1-prob)^(n-b)  
}
```

```
##iterate through different gammas (0 to 2)  
##that have different pi (gamma/(1+gamma))  
##then use lapply for summation  
##in mcneamars exact expression  
b <- 50:70  
gamma <- seq(from = 0, to = 2, by = 0.1)  
prob <- gamma/(1 + gamma)  
allresults <- list()  
for(i in 1:length(prob)){  
  results <- sapply(b, mcn.exact.p, n = 70,  
                    prob = prob[i])  
  allresults[[i]] <- results  
}  
pvals <- 2 * unlist(lapply(allresults, sum))
```

Results

- Results for original structure of table (shown on left): $\Gamma = 1$ (assume no unobserved confounding): unemployed more likely to use opioids; if unobserved confounders associated with opioids use make odds of treatment > 1.5 times higher, $p > 0.05$ so no difference

		<i>Employed</i>	
		Opioid	No opioid
<i>Unemployed</i>	Opioid	10	b = 50
	No opioid	c = 20	140

pvalues	prob	gamma
0.0000	0.4444	0.8000
0.0001	0.4737	0.9000
0.0004	0.5000	1.0000
0.0018	0.5238	1.1000
0.0057	0.5455	1.2000
0.0148	0.5652	1.3000
0.0328	0.5833	1.4000
0.0635	0.6000	1.5000

- Can also reverse table (b = employed opioid users), and see that in case of no confounding ($\Gamma = 1$), McNemar's exact p test says that employed persons *do not* use opioids at higher rate. But if there is large enough amount of unobserved confounding $\Gamma < 0.3$, null result assuming no confounding becomes significant assuming high confounding:

		<i>Unemployed</i>	
		Opioid	No opioid
<i>Employed</i>	Opioid	10	b = 20
	No opioid	c = 50	140

pvalues	prob	gamma
0.0000	0.0000	0.0000
0.0000	0.0909	0.1000
0.0178	0.1667	0.2000
0.3409	0.2308	0.3000
1.0903	0.2857	0.4000
1.6686	0.3333	0.5000
1.9089	0.3750	0.6000
1.9790	0.4118	0.7000
1.9956	0.4444	0.8000
1.9991	0.4737	0.9000
1.9998	0.5000	1.0000
2.0000	0.5238	1.1000

From Rosenbaum to Blackwell

- ▶ Takeaway from Rosenbaum:
 - ▶ The case where two persons with same observed covariates have same odds of treatment is a special case; in that special case, $\Gamma = 1$
 - ▶ We can explore the effect of deviations from that special case by seeing how our results change when two persons with the same observed covariates have different odds of treatment ($\Gamma \neq 1$)
- ▶ What does Blackwell's approach have in common? Ignorability assumption on which estimates rest is a special case; explore whether and how results change when we move away from that special case

From Rosenbaum to Blackwell

Blackwell's contributions:

1. **Reparametrize confounding function to form a more intuitive understanding of confounding's magnitude:** Draws on Imbens (2003) to compare variation in outcome explained by specific form of confounding to variation in outcome explained by particular covariates (partial R^2 for that covariate)
2. **"Evaluate alternative stories beyond one-sided bias":** while Rosenbaum bounds largely focus on one-sided bias, in theory, we can construct a bespoke confounding function well-suited to our particular confounding story; in practice, large focus on one-sided bias with some alignment bias
3. **Framing sensitivity analysis explicitly within potential outcomes framework**

Blackwell: the confounding function

- ▶ General form:

$$q(a, x) = E[Y_i(A)|A_i = a, X_i = x] - E[Y_i(A)|A_i = 1 - a, X_i = x]$$

- ▶ Single-parameter version (equation 5):

$$q(a, x; \alpha) = E[Y_i(A)|A_i = a, X_i = x] - E[Y_i(A)|A_i = 1 - a, X_i = x] = \alpha$$

- ▶ One-sided bias (equation 6):

$$q(a, x; \alpha) = \alpha(2a - 1)$$

- ▶ One-sided bias for treatment versus control:

$$\begin{aligned} q(1, x; \alpha) &= \alpha \\ q(0, x; \alpha) &= -\alpha \end{aligned}$$

- ▶ [More analytics](#)

Directions of α for examples

Using the one-sided bias function (reverse italicized for $\alpha < 0$)

Example	$\alpha > 0$; Confounders mean...
$A_i = 1 = \text{unemp.}$ $Y_i(a) = \text{opioid use}$	<i>Higher</i> potential likelihood of opioid use among those with higher prob. of unemp.
$A_i = 1 = \text{job-training (JT)}$ $Y_i(a) = \text{earnings}$	<i>Higher</i> potential earnings among those with higher prob. of participating in JT
$A_i = 1 = \text{fem. judge on panel;}$ $Y_i(a) = \text{liberal vote}$	<i>Higher</i> potential likelihood of voting liberal among those with higher probability of being in panel with a female
$A_i = 1 = \text{neg. campaign}$ $Y_i(a) = \text{turnout}$	<i>Higher</i> potential turnout among campaigns with higher prob. of going negative

Once we have the confounding function, we use it to "de-confound" each observation's observed outcome

1. Begin with each individual's observed outcome Y_i
2. Create a confounding-adjusted outcome:

$$Y_i^q = Y_i - q(A_i, X_i)Pr[1 - A_i|X_i]$$

3. Example with snapshot of LaLonde data and treatment prediction equation accounting for degree status, past earnings, age, etc. and where $\alpha = 500$ v. $\alpha = 2000$, and adjust = $q(a, x)Pr(1 - A_i|X_i)$

id	no	ai	$Pr(A_i = 1)$	$Pr(1 - A_i X_i)$	y_i	adjust $\alpha = 500$	y_i^q $\alpha = 500$	adjust $\alpha = 2000$	y_i^q $\alpha = 2000$
		deg.							
1	0	1	0.66	0.34	0	169	-169	676	-676
2	1	1	0.36	0.64	4666	319	4348	1275	3391
3	1	0	0.36	0.36	445	-181	627	-725	1170
4	0	0	0.58	0.58	12384	-289	12673	-1156	13539

Once we have the confounding function, we use it to "de-confound" each observation's observed outcome

id	no	ai	$Pr(A_i$	$Pr(1 - A_i X_i)$	y_i	adjust	y_i^q	adjust	y_i^q
	deg.		$= 1)$			$\alpha = 500$	$\alpha = 500$	$\alpha = 2000$	$\alpha = 2000$
1	0	1	0.66	0.34	0	169	-169	676	-676
2	1	1	0.36	0.64	4666	319	4348	1275	3391
3	1	0	0.36	0.36	445	-181	627	-725	1170
4	0	0	0.58	0.58	12384	-289	12673	-1156	13539

Two things happening in de-confounding of outcome variable:

- ▶ **What $q(a, x)$ is doing:** since we set $q(1, x) > q(0, x) \implies \alpha > 0$ for tx and $\alpha < 0$ for control, the earnings of those in treatment group are adjusted downwards while earnings of those in control group are adjusted upwards
 - ▶ Example of downward adjustment of treatment group: id2 goes from \$4666 to \$4348 if $\alpha = 500$ and \$3391 if $\alpha = 2000$
 - ▶ Example of upward adjustment of control group: id 3 goes from \$445 to \$627 and \$1170 when $\alpha = 500$ and 2000 respectively
- ▶ **What $Pr[1 - A_i|X_i]$ is doing:** those with higher probability of being in opposite treatment group have greater-magnitude adjustment of potential outcomes
 - ▶ Example id2: higher probability of being in *control* than tx due to no high school degree; greater downward adjustment

More intuition behind role of $Pr(1 - A_i|X)$ in adjustment

- Focus on id1 and id2:

id	no	ai	$Pr(A_i = 1)$	$Pr(1 - A_i X_i)$
	deg.			
1	0	1	0.66	0.34
2	1	1	0.36	0.64

- One (heuristic), way to think about id2s larger adjustment is to think about a person's probability of treatment being partitioned into $Pr(A_i = 1|observed_i) + Pr(A_i = 1|unobserved_i) = 1$, with id1 and id2 having different partitions (green: $Pr(A_i = 1|observed_i)$ and red: $Pr(A_i = 1|unobserved_i)$):
 - id1's partition: observed covars played larger role in fact id1 was treated

0.34	0.66
------	------
 - id2's partition: unobserved covars played larger role in fact id2 was treated (given low role for observed covars)

0.64	0.36
------	------
- For id2, we give larger downward adjustment because small role for observed covars in him being treated means we assume larger role played by unobserved covars/greater confounding

Once we have the "de-confounded" outcome, how do we proceed?

id	ai	y_i	y_i^q	y_i^q
			$\alpha = 500$	$\alpha = 2000$
1	1	0	-169	-676
2	1	4666	4348	3391
3	0	445	627	1170
4	0	12384	12673	13539

1. Switch from estimation of treatment effect with confounded outcome Y_i to estimation of treatment effect with de-confounded outcome Y_i^q , e.g., if:

► Old:

$$Y_i = \alpha + \beta_1 A_i + \beta_2 X_i + e_i$$

► New:

$$Y_i^q = \alpha + \beta_1 A_i + \beta_2 X_i + e_i$$

2. Iterate through different values in both directions (e.g., $-4000 \leq \alpha \leq 4000$) and re-estimate $\hat{\beta}_1$ (difference in earnings between treatment and control)
3. For above example, as α gets larger, the de-confounded treatment outcomes get smaller while the de-confounded control outcomes get larger, meaning at some α , $\hat{\beta}_1 = 0$

Visual illustration: full LaLonde data

```
sens.function <- function(alpha,
  data,
  outcome.mod,
  p.mod, txname,
  outcomename) {

  #add predicted probabilities of tx to data
  #and strings with log. for next step
  txwithpred <- cbind.data.frame(data, fitted(p.mod))
  tx.check <- paste(txname, "== 1", sep = " ")
  cont.check <- paste(txname, "== 0", sep = " ")

  #divide data into treatment where
  ##adjustment is alpha * Pr(1-tx)
  #and control where adjustment
  ##is -alpha * Pr(Tx) and rebind
  tx <- txwithpred %>% filter_(tx.check) %>%
  mutate(adjust = alpha * (1 - predicted_tx))
  cont <- txwithpred %>% filter_(cont.check) %>%
  mutate(adjust = -alpha * predicted_tx)
  txcont <- rbind.data.frame(tx, cont)

  #deconfound DV by subtracting adjustment
  ##and rebind
  yiq <- txcont[[outcomename]] - txcont[["adjust"]]
  txwithy <- cbind.data.frame(txcont, yiq)

  #replace dv with deconfounded dv
  new_form <- formula(outcome_model)
  new_form[[2]] <- as.name("yiq")

  ##rerun model and store tx coef and 95% CI
  adj_result <- try(lm(new_form,
    data = txwithy))
  txcoef <- coef(adj_result)[txname]
  ci <- confint(adj_result)[txname, ]
  all_est[["alpha"]] <- c(txcoef, ci)
}
```

Produces same results as *causalsens* with confound = "one.sided"

Visual illustration: full LaLonde data




Figure: ATT on left corresponds to Figure 1 in paper; ATE is on right. Difference stems from ATE: adjust both treatment and control y_i versus ATT: adjust only control y_i with $-\alpha \Pr(A_i = 1|X_i)$

Summing up thus far and where we're going next

- ▶ We reviewed the basic confounding function
- ▶ We illustrated how to use that function to 'de-confound' the outcome variable and re-estimate the treatment effect
- ▶ Now, we'll discuss three extensions of basic case:
 1. Change confounding function from one-sided bias to alignment bias
 2. Re-parametrize confounding function to express magnitude of confounding in terms of variation in R^2 rather than difference in mean inherent potential outcomes between treatment and control (α)
 3. Going from static, cross-sectional treatment assignment to treatments over time (dynamic case)

Extension one: change confounding function

- ▶ In previous example, checked how results change due to *one-sided bias* ($q(1, x) = \alpha$; $q(0, x) = -\alpha$), which captures the situation where we expect those who select into treatment to have inherently higher or lower values of outcome than those who select into control
 - ▶ Lalonde example: those who opt for job training have inherently higher or lower potential earnings than those who do not opt into treatment
- ▶ Different confounding model: *alignment bias* ($q(1, x) = q(0, x) = \alpha$) captures the situation where we expect those who select into treatment to have larger treatment effects than those who select into control
 - ▶ Lalonde example: those who opt for job training have unobserved characteristics that help them benefit more from job training than if those who did not opt for job training were subject to the treatment



Extension one: alignment bias with same LaLonde participants

id	no	ai	$Pr(A_i = 1)$	$Pr(1 - A_i X_i)$	y_i	adjust $\alpha = 500$	y_i^q $\alpha = 500$	adjust $\alpha = 2000$	y_i^q $\alpha = 2000$
1	0	1	0.66	0.34	0	169	-169	676	-676
2	1	1	0.36	0.64	4666	319	4347	1275	3391
3	1	0	0.36	0.36	446	181	265	725	-279
4	0	0	0.58	0.58	12384	289	12095	1156	11228

What changed from one-sided bias de-confounding?

1. Rather than adjusting earnings of treatment units downwards and earnings of control units upwards, all receive same sign of adjustment (downward when $\alpha > 0$)
2. Because sign of adjustment is same between treatment and control, what leads to $\hat{\beta} \rightarrow 0$ (might) be something like the following case:
 - 2.1 People with unobserved characteristics that lead them to benefit most from treatment actually have lower $Pr(A_i = 1 | X_i)$
 - 2.2 Lower $Pr(A_i = 1 | X_i) \implies$ higher $Pr(1 - A_i | X_i) \implies$ greater downward adjustment of treatment group's outcomes
 - 2.3 This is happening less (smaller downward adjustment) in controls

ATT: alignment bias (left) versus one-sided bias (right) with LaLonde data



Extension two: express magnitude of confounding in terms of R^2 rather than in terms of α

- ▶ In previous example, we measured magnitude of confounding in terms of difference between mean potential outcomes for treatment group and control (e.g., \$500 difference in inherent potential earnings, \$2000 difference in potential earnings)
- ▶ Say positive treatment effect is still significant at $\alpha = 500$ but no longer significant at $\alpha = 2000$, difficult to know which is a more plausible magnitude of confounding
- ▶ More intuitive way to measure magnitude: instead of incrementing through α , increment through different proportions of variance in outcome explained by selection into treatment (or other confounding process)
- ▶ Can then compare to partial variance explained by influential covariates
 - ▶ More robust results: treatment effect still holds even when, for instance, outcome variance explained by confounding is larger than outcome variance explained by influential observed covar.

Extension two: mechanics

1. Start with proportion of potential outcome variance explained by observed covariates X and treatment status A :

$$R_q^2(X_i, A_i) = 1 - \frac{\text{var}[Y_i(0)|X_i, A_i, q]}{\text{var}[Y_i(0)]}$$

2. Then, find the proportion of potential outcome variance explained by observed covariates X :

$$R_q^2(X_i) = 1 - \frac{\text{var}[Y_i(0)|X_i, q]}{\text{var}[Y_i(0)]}$$

3. We can express the proportion of unexplained variance in the potential outcomes due to selection into treatment by taking the variance explained by X and selection into treatment and subtracting out the variance explained by X ($R_q^2(X_i, A_i) - R_q^2(X_i)$) and rearranging to get:

$$R_q^2(A_i) = 1 - \frac{\text{var}[Y_i(0)|X_i, A_i, q]}{\text{var}[Y_i(0)|X_i, q]}$$

Extension two: mechanics

- ▶ Another way to express (3) from previous slide is as:

$$R_q^2(A_i) = 1 - \frac{\text{var}[Y_i(0)|X_i, A_i, q]}{\text{var}[Y_i(0)|X_i, q]} = 1 - \frac{\text{unrestricted model}}{\text{restricted model}} = 1 - \frac{\text{var}(e_i)}{\text{var}(e'_i)}$$

- ▶ Then, assuming q is expressed as one-sided bias:

1. Start with:

$$R_q^2(A_i) = 1 - \frac{\text{var}(e_i)}{\text{var}(e'_i)}$$

2. Simplify and plug in $e'_i = \alpha A_i + e_i$ to numerator:

$$R_q^2(A_i) = \frac{\text{var}(\alpha A_i + e_i) - \text{var}(e_i)}{\text{var}(e'_i)}$$

3. Use $\text{var}(aX) = a^2 \text{var}(x)$ and $\text{var}(x + y) = \text{var}(x) + \text{var}(y)$ to further simplify:

$$R_q^2(A_i) = \frac{\alpha^2 \text{var}(A_i) + \text{var}(e_i) - \text{var}(e_i)}{\text{var}(e'_i)} = \frac{\alpha^2 \text{var}(A_i)}{\text{var}(e'_i)}$$

Extension two: mechanics

- ▶ α^2 and $\text{var}(A_i)$ are straightforward to estimate, what about $\text{var}(e'_i)$?
- ▶ Remember that $\text{var}[Y_i(0)|X_i, q] = \text{var}(e'_i)$ corresponds to the model where $A_i = 0$, so:

$$Y_i(0) = X_i\beta + \alpha A_i + e_i$$


becomes the restricted model:

$$Y_i(0) = X_i\beta + e'_i$$

- ▶ Because assuming the confounding function is correct, $E[Y_i(0)] = E[Y_i^q]$, we can estimate $Y_i(0) = X_i\beta + e'_i$ by:
 1. Regressing Y_i^q on X for $A_i = 0$
 2. Finding variance of residuals
 3. Can then find partial R^2 for particular covariates of interest to which to compare this value²

²Implementation-wise, can use drop1 in R to find sum of squares and residual sum of squares when you restrict the model to a particular variable

Extension two: graph for LaLonde ATT and ATE, also plot partial R^2 of education ($p < 0.1$ in original outcome model)



Note when implementing: since a bit unclear what negative R^2 is, plots in paper seem to be generated by iterating through positive R^2 values, splitting the data at $\alpha = 0$, and then setting R^2 to $-R^2$ for observations where $\alpha < 0$

Extension three: time-varying treatments and confounding

Motivation:

- ▶ Previous examples were treatments and covariates observed at one point in time (e.g., one-shot job training)
- ▶ Blackwell (2012) argues that at least some treatments social scientists care about are composed of *action histories*, a specific sequence of treatments
- ▶ Example: decisions to run negative ($A_i = 1$) versus positive campaign ads at different weeks leading up to an election

	campaign	Week 1	Week 2	Week 3	Week 4	Vote share
	1	Neg	Neg	Neg	Pos	67%
	2	Pos	Neg	Pos	Neg	30%
	3	Neg	Neg	Neg	Neg	47%

- ▶ These cases have more thorny dilemma than the single-shot treatment confounder issue: *time-varying confounders*, which are both affected by past treatments and influence choice of treatment at time t (e.g., poll results from week 2 being influenced by neg v. pos ad at week 1 and influencing probability of negative campaign at week 3)

Extension three: from ignorability to sequential ignorability

- ▶ Single-shot treatment case: rests on ignorability assumption and sensitivity analysis probes how results change with violations
- ▶ Dynamic treatment case: rests on sequential ignorability assumption; likewise, sensitivity analysis probes consequences of violations (notation: history of a variable up to time t):

$$Y_i(\underline{a}) \perp A_{it} | \underline{X}_{it}, \underline{A}_{it-1}$$

- ▶ Confounding function in dynamic treatment case:

$$\begin{aligned} q_t(\underline{a}, \underline{x}_t) &= E[Y(\underline{a}) | A_t = a_t, \underline{X}_t = \underline{x}_t, \underline{A}_{t-1} = \underline{a}_{t-1}] \\ &\quad - E[Y(\underline{a}) | A_t = 1 - a_t, \underline{X}_t = \underline{x}_t, \underline{A}_{t-1} = \underline{a}_{t-1}] \end{aligned}$$

- ▶ How do we use that confounding function to adjust the outcome variable?

$$Y_i^\alpha = Y_i - \sum_{t=0}^T q_t(\underline{A}_{it}, \underline{X}_{it}; \alpha) Pr(A_t = 1 - A_{it} | \underline{A}_{it-1}, \underline{X}_{it})$$



Extension three mechanics: de-confounding outcome in dynamic case

1. First, model the probability of treatment in week t conditional upon $\underline{X}_{it}$ and $\underline{A}_{it}$. E.g., letting $A_{it} = 1$ = negative campaign ad for campaign i in week t if $t = 3$:

$$\begin{aligned} Pr(A_{i3} = 1) = & \alpha + \beta_1 neg_{i1}(1 = \text{yes}; 0 = \text{no}) + \beta_2 neg_{i2}(1 = \text{yes}; 0 = \text{no}) \\ & + \beta_3 donation_{i1} + \beta_4 donation_{i2} + \beta_4 donation_{i3} + e_i \end{aligned}$$

2. Using fitted values from 1, assign each i at each t the following, which is 1 minus the probability of reaching this treatment history: $Pr(A_t = 1 - A_{it} | \underline{A}_{it-1}, \underline{X}_{it})$ and multiply by appropriate α (e.g, alignment versus one-sided bias)
3. For each i , sum the results of (2) across t and subtract from Y_i to create Y_i^α , which is then used in whatever estimation procedure for treatment effect is chosen (e.g., marginal structural model)

Extension three: Blackwell's results for negative campaign case, discuss interpretation



Briefly: alternative approaches (Morgan and Winship)

- ▶ Blackwell and Rosenbaum correspond to Section 12.3 ("Sensitivity Analysis for Provisional Causal Effects Estimate")
- ▶ Authors outline a different approach in Section 12.2 where rather than investigating the sensitivity of a treatment's *point estimate* to violations of *strong* assumptions like ignorability, we place bounds on the causal effect of treatment by adding *weak assumptions*
- ▶ Bounding, with work by Manski and more recently Keele, outlines different strategies for bounding causal quantities of interest like the Average Treatment Effect (ATE):
 1. *No-assumptions bounds*: if outcome variable is constrained to lie between 0 and 1, the initial bounds on the treatment effect $ATE = [-1, 1]$ (width = 2) can be shrunk to bounds where width = 1 by assuming different combinations of quantities for unobserved outcomes (e.g., $E[Y = 1|A = 0] = 1$ and $E[Y = 0|A = 1] = 0$, and vice versa)
 2. *Weak assumptions bounds*: can narrow width of interval (most helpfully so that the bounds exclude 0) through various weak assumptions: e.g., monotone treatment response, monotone treatment selection

Recap/some questions

1. What we reviewed: why sensitivity analysis? Rosenbaum bounds approach; Blackwell approach: simple case and three extensions (different confounding function; re-parametrization of magnitude of confounding; dynamic treatment)...some questions:
2. One advantage of Blackwell is flexibility to create a confounding function specific to a theoretical story. Yet most examples drew on one-sided bias (with some alignment bias). Other ideas for confounding functions beyond these?
3. Blackwell argues that one limitation of his approach is that at its core, it relies on a "selection on the observables" assumption so it is "incompatible with certain other approaches to causal inference such as instrumental variables"(p. 181). Yet Morgan and Winship argue that "selection-bias models are most effectively estimated" when we include instrumental variables in Z . For the step in Blackwell's approach where we estimate each unit's probability of treatment, can we use an instrumental variable as part of this estimation?
4. Dynamic treatment approach (action history) was developed in biostatistics where we care about quantities like cumulative treatment history. We have Blackwell's negative ad example; what other social science questions would benefit from a treatment-history type approach (with accompanying sensitivity checks)?

Appendix

More analytics behind de-confounding outcome

1. Begin by expanding $E[Y_i(0)]$ in case without covariates:

$$E[Y_i(0)] = E[Y_i(0)|A_i = 0]Pr[A_i = 0] + E[Y_i(0)|A_i = 1]Pr[A_i = 1]$$

2. Since the terms in red below sum to zero, we can add them to equation, with undersets used in next step:

$$E[Y_i(0)] = \underset{A}{E[Y_i(0)|A_i = 0]Pr[A_i = 0]} + \underset{B}{E[Y_i(0)|A_i = 1]Pr[A_i = 1]} +$$

$$\underset{C}{E[Y_i(0)|A_i = 0]Pr[A_i = 1]} - \underset{D}{E[Y_i(0)|A_i = 0]Pr[A_i = 1]}$$

3. Rewrite (A) as $E[Y_i(0)|A_i = 0]Pr[A_i = 0] = E[Y_i(0)|A_i = 0](1 - Pr[A_i = 1]) = E[Y_i(0)|A_i = 0] - E[Y_i(0)|A_i = 0]Pr[A_i = 1]$ and substitute into above:

$$E[Y_i(0)] = E[Y_i(0)|A_i = 0] - \tag{1}$$

$$E[Y_i(0)|A_i = 0]Pr[A_i = 1] + \tag{2}$$

$$E[Y_i(0)|A_i = 1]Pr[A_i = 1] + \tag{3}$$

$$E[Y_i(0)|A_i = 0]Pr[A_i = 1] - \tag{4}$$

$$E[Y_i(0)|A_i = 0]Pr[A_i = 1] \tag{5}$$

4. Cancel out (2) and (5), combine (3) and (4) and change sign by subtracting negative, and multiply (1) by $Pr[A_i = 0] + Pr[A_i = 1] = 1$ (typo in text) to get **red** = $q(0)$, **cyan** = $Y_i(0)$, and **green** = $Pr(1 - A_i)$:

$$E[Y_i(0)] = \text{cyan}[E[Y_i(0)|A_i = 0](Pr[A_i = 0] + Pr[A_i = 1])$$

$$- \text{green}[Pr[A_i = 1](E[Y_i(0)|A_i = 0] - E[Y_i(0)|A_i = 1])]$$