Causal Inference and Experimentation

Macartan Humphreys — mh2245@columbia.edu

November 15, 2011

Experiments

- Experiments are investigations in which an intervention, in all its essential elements, is under the control of the investigator. Cox Reid (2000) define
- Two major types of control:
  1. over assignment to treatment – this is at the heart of many field experiments
  2. control over the treatment itself – this is at the heart of many lab experiments
- Both important. Main focus today is on 1 and on the question: *how does control over assignment treatment allow you to make reasonable statements about causal effects?*

Key ideas

The big ideas:

- The potential outcomes framework: How you can think about causality without functional forms
- How random assignment to treatment is actually random sampling from alternative universes
- Randomization: How to do it
- Analysis: Why you should stop running regressions
- Randomization inference: How you can exploit randomization for statistical tests without any assumptions about distributions
- Analysis: LATE What you are really estimating in an encouragement design
- How to think about spillovers
- What this is and isn't good for

Motivation

- Say you want to know if a particular intervention (like aid) caused a particular outcome (like good governance) you need to know:
  1. What is the outcome?
  2. What would the outcome have been if there were no intervention?
- The problem
  1. . . . this is hard
  2. . . . this is impossible

  The problem in 2 is that you need to know what would have happened if things were different. You need information on a **counterfactual**

The Fundamental Problem of Causal Inference

- The best we can do is to make a comparison.
- **Problem**: With what units can we make a meaningful comparison?
- Illustration:
  - **Question:** Do UN peacekeeping missions actually bring about peace?
  - First Evidence: No if you compare outcomes in areas where UN Peacekeepers work to outcomes in areas where UN Peacekeepers do not work you will see that there is less security in places where they do not work.

The Fundamental Problem of Causal Inference

- The best we can do is to make a comparison.
- **Problem:** With what units can we make a meaningful comparison?
- Illustration:
  - **Question:** Do UN peacekeeping missions actually bring about peace?
  - First Evidence: Just compare outcomes in Congo with outcomes in Kitsilano.

    The UN is a disaster!

The Fundamental Problem of Causal Inference

- **Problem**: Comparing outcomes in places with and without treatment only makes sense if the areas you compare are comparable.
- In fact the UN tends to go to hard places and thats why things look so bad when you do a simple comparison.
- So the right answer might be Yes!
- For comparisons to be valid, outcomes in comparison units have to look like what outcomes *would* have looked like in treatment communities.

The Potential Outcomes Framework

- For each unit (say community) we assume that there are two post-intervention outcomes: $Y_i(1)$ and $Y_i(0)$.
- eg $Y(1)$ is the outcome that **would** obtain *if* the unit received the treatment.
- The **causal effect** of Treatment (relative to Control) is:

$$\tau_i = Y_i(1) - Y_i(0)$$

- Note:
  - the causal effect is defined at the *individual level*.
  - there is no "data generating process" or functional form
  - the causal effect is defined relative to something else and so a counterfactual must be conceivable (did Germany cause the second world war?)
  - are there any substantive assumptions made here so far?

The Potential Outcomes Framework

- What *do* we observe?
- Say Z_i indicates whether the unit i is assigned to treatment ($Z_i = 1$) or not ($Z_i = 0$) . It describes the treatment process. Then what we observe is:

$$Y_i = Z_i Y_i(1) + (1 - Z_i) Y_i(0)$$

- Say Z is a random variable, then this is a sort of data generating process. BUT the key things to note is
  - Y_i is random but the randomness comes from Z_i — the potential outcomes, $Y_i(1)$, $Y_i(0)$ are fixed
  - Compare this to a regression approach in which Y is random but the X's are fixed. eg:

$$Y \sim N(\beta X, \sigma^2) \text{ or } Y = \alpha + \beta X + \epsilon, \epsilon \sim N(0, \sigma^2)$$

The Potential Outcomes Framework

- The causal effect of Treatment (relative to Control) is:

$$\tau_i = Y_i(1) - Y_i(0)$$

- This is what we want to estimate
- BUT: We never can observe both $Y_i(1)$ and $Y_i(0)$!
- This is the **fundamental problem** (Holland)

The Potential Outcomes Framework

- Now for some magic. We really want to estimate:

$$\tau_i = Y_i(1) - Y_i(0)$$

- BUT: We never can observe both $Y_i(1)$ and $Y_i(0)$!
- Say we lower our sights and try to estimate an average treatment effect:

$$\tau = E(Y(1) - Y(0))$$

- Now make use of the fact that

$$E(Y(1) - Y(0)) = E(Y(1)) - E(Y(0))$$

- In words: *The average of differences is equal to the difference of averages*; here, the average treatment effect is equal to the difference in average outcomes in treatment and control units.
- The magic is that *while we can't hope to measure the differences; we are good at measuring averages*.

The Potential Outcomes Framework

- So we want to estimate $E(Y(1))$ and $E(Y(0))$.
- We know that we can estimate averages of a quantity by taking the average value from a random sample of units
- To do this here we need to select a random sample of the $Y(1)$ values and a random sample of the $Y(0)$ values, in other words, we **randomly assign** subjects to treatment and control conditions.
- When we do that we can in fact estimate:

$$E_N(Y_i(1)|Z_i = 1) - E_N(Y_i(0)|Z_i = 0)$$

which in expectation equals:

$$E(Y_i(1)|Z_i = 1 \text{ or } Z_i = 0) - E(Y_i(0)|Z_i = 1 \text{ or } Z_i = 0)$$

- This highlights a deep connection between random assignment and random sampling: when we do random assignment *we are in fact randomly sampling from different possible worlds*.

The Potential Outcomes Framework

- It also provides a positive argument for causal inference from randomization, rather than simply saying with randomization "everything else is controlled for"
- Where are the covariates?

The Potential Outcomes Framework

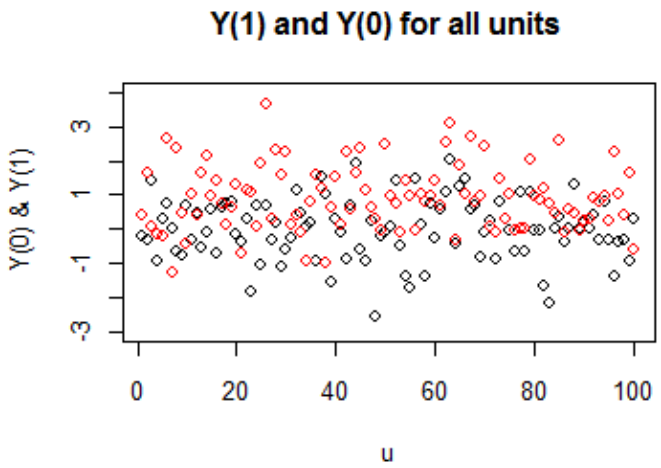

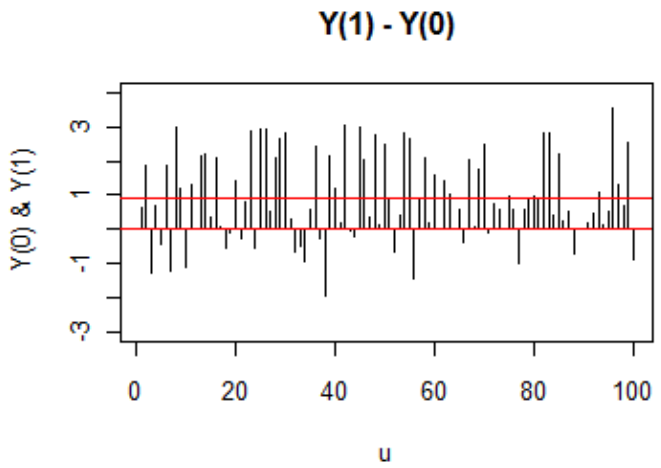

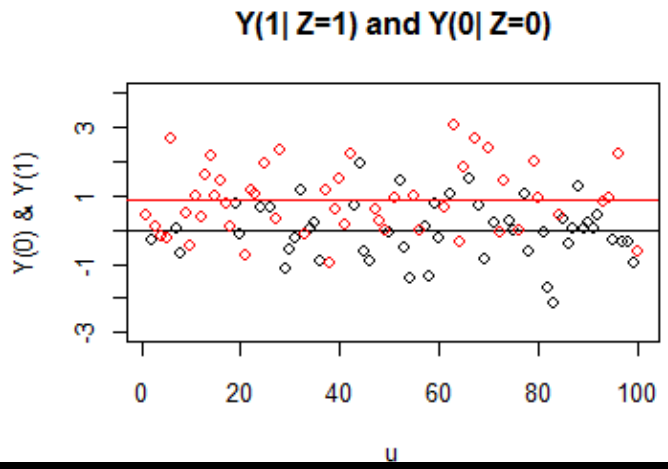

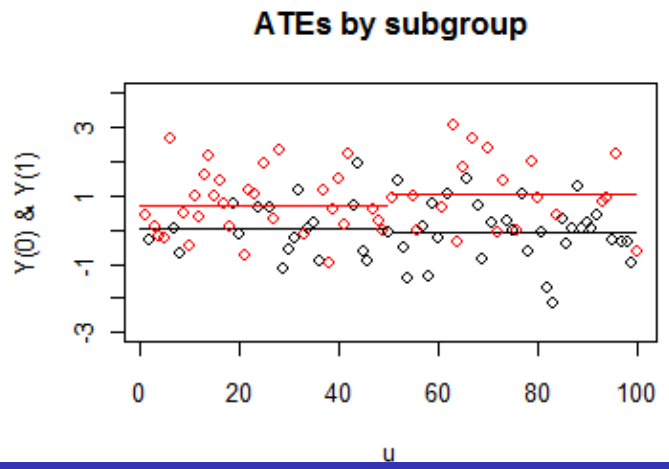

The Potential Outcomes Framework: Covariates?

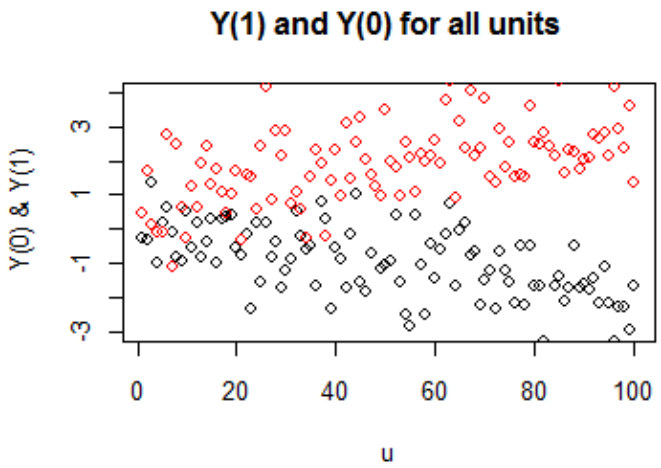

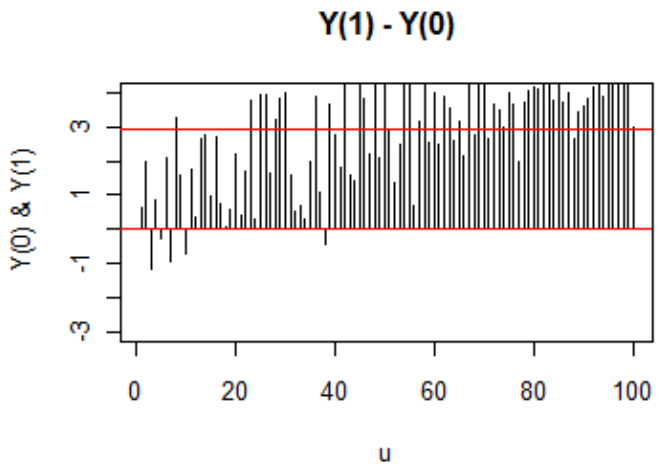

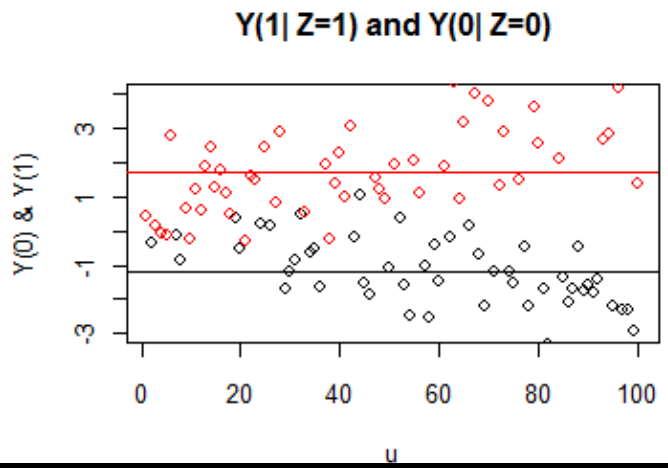

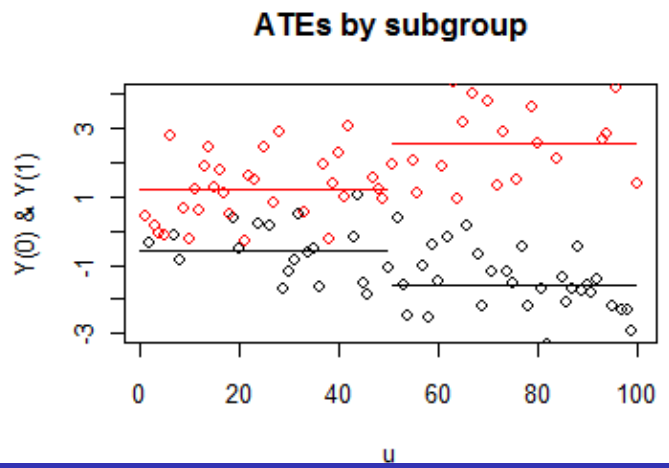

Code for potential outcomes graphs

```
par(mfrow=c(2,2))
N=100; u = seq(1:N);  Y0 = rnorm(N);  Y1 = rnorm(N) + 1; Z= 1:N %in% sample(N, N/2)
po.graph = function(N, Y0,Y1,u, Z){
yl = "Y(0) & Y(1)"

plot(u, Y0, ylim=c(-3, 4), xlim=c(1,N), xlab="u", ylab=yl)
   lines(u, Y1, type = "p", col="red")
 title("Y(1) and Y(0) for all units ")
plot(u, Y1-Y0, type = "h", ylim=c(-3, 4), xlim=c(1,N),main = "Y(1) - Y(0)", xlab="u", ylab=yl)
 abline(h=0, col="red"); abline(h=mean(Y1-Y0), col="red")
plot(u[Z==0], Y0[Z==0], ylim=c(-3, 4), xlim=c(1,N), main = "Y(1| Z=1) and Y(0| Z=0)", xlab="u", ylab=yl)
 abline(h=mean(Y0[Z==0]))
 lines(u[Z==1], Y1[Z==1], type = "p", col="red")
 abline(h=mean(Y1[Z==1]), col="red")
plot(u[Z==0&u<=N/2], Y0[Z==0&u<=N/2], ylim=c(-3, 4), xlim=c(1,N), main = "Subgroup ATEs", xlab="u", ylab =
 segments(0, mean(Y0[Z==0 & u<=N/2]), N/2, mean(Y0[Z==0 & u<=N/2]), lwd =  1.3)
 lines(u[Z==1 & u<=N/2], Y1[Z==1 & u<=N/2], type="p",ylim=c(-3, 4), col="red")
 segments(0, mean(Y1[Z==1 & u<=N/2]), N/2, mean(Y1[Z==1 & u<=N/2]), lwd =  1.3, col="red")
 lines(u[Z==0 & u>N/2], Y0[Z==0 & u>N/2], type = "p", ylim=c(-3, 4), xlim=c(1,N))
 segments(1+N/2, mean(Y0[Z==0 & u>N/2]), N, mean(Y0[Z==0 & u>N/2]), lwd =  1.3)
 points(u[Z==1 & u>N/2], Y1[Z==1 & u>N/2], type="p", ylim=c(-3, 4), col="red")
 segments(1+N/2, mean(Y1[Z==1 & u>N/2]), N, mean(Y1[Z==1 & u>N/2]), lwd =  1.3, col="red")
}
po.graph(N, Y0,Y1,u,Z)
po.graph(N, Y0-u/50,Y1+u/50,u,Z)
```

Estimands: ATE, ATT, ATC, S-, P-, C-, ITT, LATE

The key estimands and estimators are:

$$
\begin{aligned}
\tau_{ATE} \equiv \quad & \mathrm{E}(\tau_i) && = \quad \sum_x \frac{w_x}{\sum_j w_j}\tau_x && \widehat{\tau}_{ATE} = \quad \sum_x \frac{w_x}{\sum_j w_j}\widehat{\tau}_x \\
\tau_{ATT} \equiv \quad & \mathrm{E}(\tau_i | \, Z_i = 1) && = \quad \sum_x \frac{p_x w_x}{\sum_j p_j w_j}\tau_x && \widehat{\tau}_{ATT} = \quad \sum_x \frac{p_x w_x}{\sum_j p_j w_j}\widehat{\tau}_x \\
\tau_{ATC} \equiv \quad & \mathrm{E}(\tau_i | \, Z_i = 0) && = \quad \sum_x \frac{(1-p_x) w_x}{\sum_j (1-p_j) w_j}\tau_x && \widehat{\tau}_{ATC} = \quad \sum_x \frac{(1-p_x) w_x}{\sum_j (1-p_j) w_j}\widehat{\tau}_x
\end{aligned}
$$

where x indexes strata, p_x is the share of units in each stratum that is treated, and w_x is the size of a stratum.
Here:

- ATE is Average Treatment Effect (all units)
- ATT is Average Treatment Effect on the Treated
- ATC is Average Treatment Effect on the Controls

Estimands: ATE, ATT, ATC, S-, P-, C-, ITT, LATE

In addition, each of these can be calculated:

- for the **population**, in which case we refer to PATE, PATT, PATC and $\widehat{PATE}, \widehat{PATT}\,\widehat{PATC}$
- for a **sample**, in which case we refer to SATE, SATT, SATC, and $\widehat{SATE}, \widehat{SATT}\,\widehat{SATC}$

And for different subgroups,

- given some value on a covariate, in which case we refer to CATE (conditional average treatment effect)
- for unobservable subgroups, we estimate LATE (Local Average Treatment Effect (see below).

With non-compliance we might estimate ITT —the "intention to treat" effect

Basic randomization

Design: How to Randomize

Basic randomization

- Basic randomization is very simple. For example, say you want to assign 5 of 10 units to treatment in R.

```
1:10 %in% sample(1:10, 5)
 [1]  TRUE FALSE  TRUE FALSE FALSE  TRUE  TRUE FALSE FALSE  TRUE
```

- But in general you might want to set things up so that your randomization is **replicable**. You can do this by setting a **seed**:

```
set.seed(20111112)
1:10 %in% sample(1:10, 5)
 [1]  TRUE  TRUE FALSE FALSE  TRUE FALSE FALSE  TRUE FALSE  TRUE
set.seed(20111112)
1:10 %in% sample(1:10, 5)
 [1]  TRUE  TRUE FALSE FALSE  TRUE FALSE FALSE  TRUE FALSE  TRUE
```

Basic randomization

- Even better is to set it up so that it can reproduce **lots of possible draws** so that you can check the propensities for each unit.

```
set.seed(20111112)
P = sapply(1:1000, function(i) 1:10 %in% sample(1:10, 5))
apply(P, 1, mean)
 [1] 0.525 0.486 0.502 0.500 0.511 0.491 0.485 0.484 0.501 0.515
```

- Here the P matrix gives 1000 possible ways of allocating 5 of 10 units to treatment. We can then confirm that the average propensity is 0.05.
- A huge advantage of this approach is that if you make a mess of the random assignment; **you can still generate the P matrix and use that for all analyses**!

Basic randomization: Fixer

Say you made a mess and used a randomization that was correlated with some variable, X. For example:

- The randomization is done in a way that introduces a correlation between Treatment Assignment and Potential Outcomes
- Then possibly, even though there is no true causal effect we naively estimate a large one — enormous bias
- However since we know the assignment procedure we can **fully** correct for the bias
- In the next example we do this using "inverse propensity score weighting." This is exactly analogous to standard survey weighting — since we selected different units for treatment with different probabilities, we weight them differently to recover the average outcome among treated units (same for control).

Basic randomization: Fixer

Say you made a mess and used a randomization that was correlated with some variable, X. Then you can still use information on the assignment process to recover the right estimates.

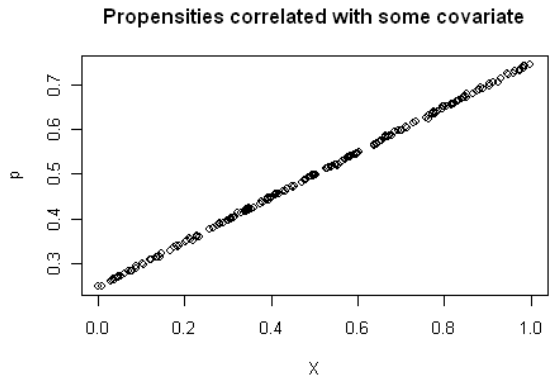

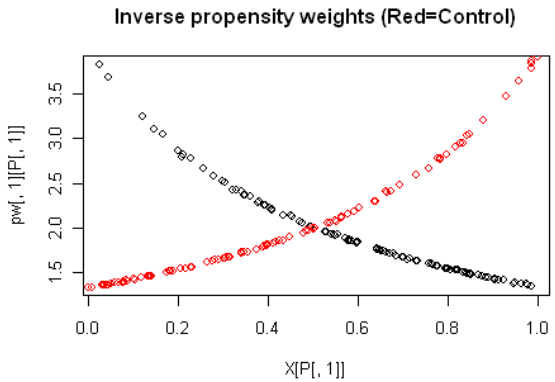

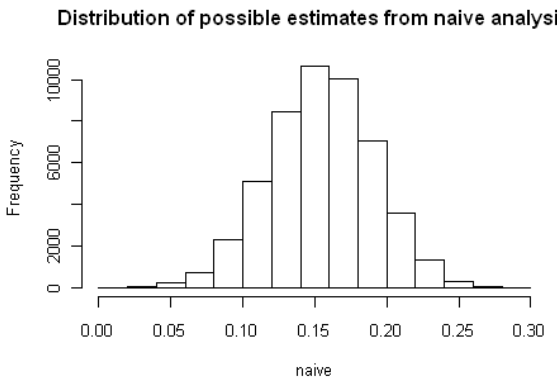

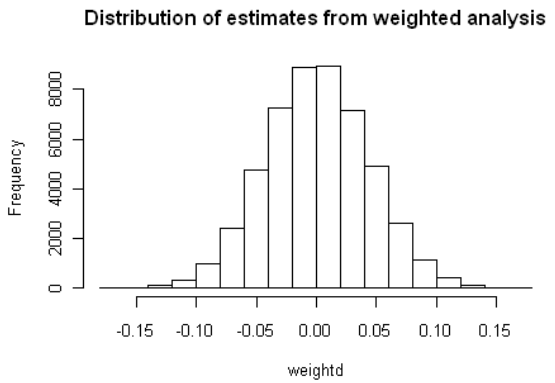

Basic randomization: Fixer

Code for these graphs:

```
n = 200; reps = 50000; X = runif(n)      # Create a covariate (length n)
Y <-Y1<-Y0<-X                            # Say X completely determines Y!
Z = function(i) rank(X+2*runif(n))>(n/2) # Bad randomization!
P = sapply(1:reps, Z)                    # Lots of possible draws
p = apply(P, 1, mean)                    # Recreate propensities!
pw = (!P)*(1/(1-p)); pw[P]=(P*(1/p))[P]  # Create inv prop weights
naive  =sapply(1:ncol(P),function(i) {mean(Y[P[,i]])-mean(Y[!P[,i]])})
weightd=sapply(1:ncol(P),function(i) {weighted.mean(Y[P[,i]], pw[,i][P[,i]])-
   weighted.mean(Y[!P[,i]], pw[,i][!P[,i]])}) # IPW estimates
par(mfrow=c(2,2));
   plot(X, p, main="Propensities correlated with some covariate");
   plot(X[P[,1]], pw[,1][P[,1]],
   main="Inverse propensity weights (Red=Control)");
   points(X[!P[,1]], pw[,1][!P[,1]], col="red")
   hist(naive, main="Distribution of possible estimates from naive analysis");
   hist(weightd, main="Distribution of estimates from weighted analysis")
```

Blocking

There are more or less **efficient** ways to randomize.

- Randomization helps ensure good balance on all covariates (observed and unobserved) *in expectation*.
- But balance may not be so great *in realization*
- Blocking can help ensure balance ex post on observables

Consider a case with 4 units and two strata. There are 6 possible assignments of 2 units to treatment:

ID	X	Y(0)	Y(1)	R1	R2	R3	R4	R5	R6
1	1	0	1	1	1	1	0	0	0
2	1	0	1	1	0	0	1	1	0
3	2	1	2	0	1	0	1	0	1
4	2	1	2	0	0	1	0	1	1
$\widehat{\tau}$:				0	1	1	1	1	2

Even with a constant treatment effect and everything uniform within blocks, there is variance in the estimation of $\widehat{\tau}$. This can be elimination by excluding R1 and R6.

Blocking

- Blocking is a case of **restricted randomization**. Although each unit is sampled with equal probability, the *profiles* of possible assignments are not.
- You have to take account of this when doing analysis
- There are many other approaches.
  - "**Matched Pairs**" are a particularly fine approach to blocking
  - You could also randomize and then **replace the randomization** if you do not like the balance. This sounds tricky (and it is) but it is OK as long as you understand the true lottery process you are employing and incorporate that into analysis
  - It is even possible to block on **covariates for which you don't have data** ex ante, by using methods in which you allocate treatment over time as a function of features of your sample (also tricky)

Blocking

Simple blocking in R (5 pairs):

```
> sapply(1:5, function(i) rank(runif(2))==1)
      [,1]  [,2]  [,3]  [,4]  [,5]
[1,]  TRUE FALSE  TRUE FALSE FALSE
[2,] FALSE  TRUE FALSE  TRUE  TRUE
```

Factorial Designs

- Often when you set up an experiment you want to look at more than one treatment.
- Should you do this or not? How should you use your power?

	$T2 = 0$	$T2 = 1$
T1=0	50%	0%
T1= 1	50%	0%

	$T2 = 0$	$T2 = 1$
T1= 0	25%	25%
T1 = 1	25%	25%

	$T2 = 0$	$T2 = 1$
T1= 0	33.3%	33.3%
T1= 1	33.3%	0%

Factorial Designs

- Surprisingly adding multiple treatments does not eat into your power
- Especially when you use a fully crossed design like the middle one above.
- Fisher: "No aphorism is more frequently repeated in connection with field trials, than that we must ask Nature few questions, or, ideally, one question, at a time. The writer is convinced that this view is wholly mistaken."
- However – adding multiple treatments *does* alter the **interpretation** of your treatment effects. If T2 is an unusual treatment for example, then half the T1 effect is measured for unusual situations.

Factorial Designs: In practice

- In practice if you have a lot of treatments it can be hard to do full factorial designs – there may be too many combinations.
- In such cases people use **fractional factorial designs**, like the one below (5 treatments but only 8 units!)

Variation	T1	T2	T3	T4	T5
1	0	0	0	1	1
2	0	0	1	0	0
3	0	1	0	0	1
4	0	1	1	1	0
5	1	0	0	1	0
6	1	0	1	0	1
7	1	1	0	0	0
8	1	1	1	1	1

- In R, look at `library(survey); hadamard(7)`

Covariate Adjustment

Analysis

ATE and Var(ATE)

Unbiased estimates of the (sample) average treatment effect can be estimated (even if there is imbalance on covariates!) using:

$$\widehat{ATE} = \frac{1}{n_T}\sum_T Y_i - \frac{1}{n_C}\sum_C Y_i,$$

You can also estimate variance straight from the data. Again, conditioning on the sample, we have:

$$V(\widehat{ATE}) = \frac{n}{n-1}\left[\frac{V(Y(1))}{n_T} + \frac{V(Y(0))}{n_C}\right] - \frac{1}{n-1}\left[V(Y(1)) + V(Y(0)) - 2C\left(Y(1), Y(0)\right)\right]$$

- . . . where V denotes variance and C covariance
- Use sample estimates $s^2(\{Y_i\}_{i=1}^{M})$ and $s^2(\{Y_i\}_{i=M+1}^{N})$ for the first part.
- $C(Y(1), Y(0))$ cannot be estimated from data.
- The "Neyman" estimator ignores the second part (and so is conservative).

Illustration of Neyman Conservative Estimator

Illustration of how conservative the conservative estimator of variance really is (numbers in plot are correlations between $Y(1)$ and (0). Here we confrim that (i) the estimator is conservative (ii) the estimator is more conservative for negative correlations between $Y(0)$ and $Y(1)$ — eg if those cases that do particularly badly in control are the ones that do particularly well in treatment

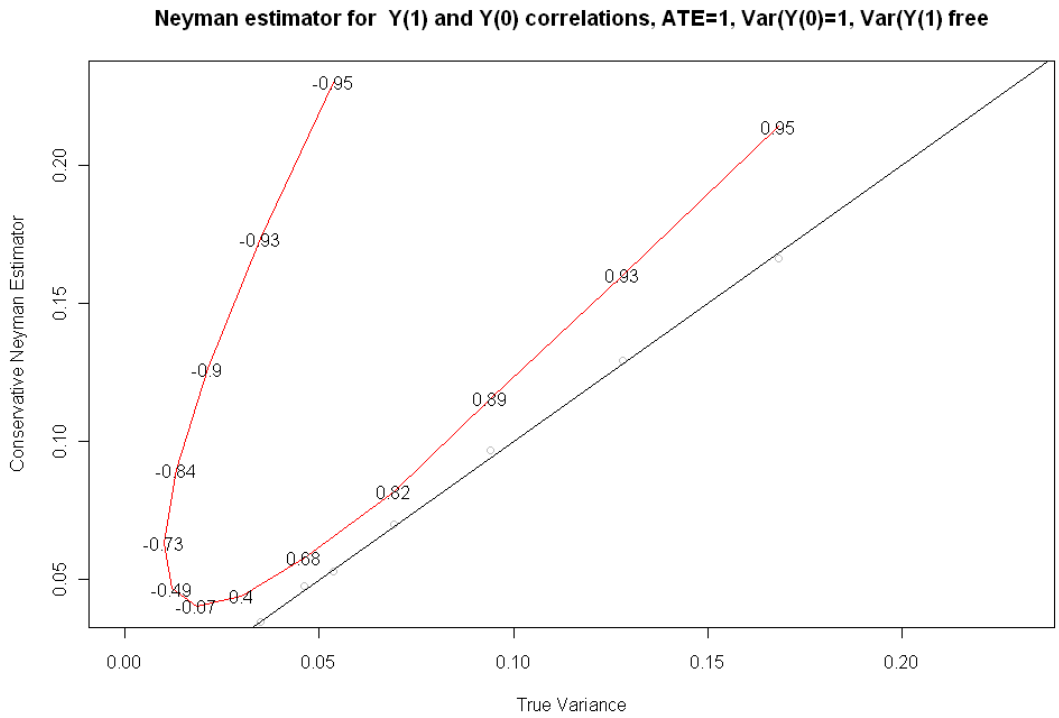

Illustration of Neyman Conservative Estimator

τ	ρ	$\sigma^2_{Y(1)}$	Δ	σ^2_τ	$\widehat{\sigma}^2_\tau$	$\widehat{\sigma}^2_{\tau(\text{Neyman})}$
1	-0.951	10.398	-0.177	0.054	0.053	0.230
1	-0.932	7.581	-0.139	0.035	0.035	0.174
1	-0.901	5.265	-0.105	0.021	0.021	0.126
1	-0.843	3.449	-0.077	0.013	0.013	0.090
1	-0.730	2.133	-0.053	0.010	0.010	0.063
1	-0.494	1.316	-0.035	0.012	0.012	0.047
1	-0.066	1.000	-0.022	0.018	0.019	0.040
1	0.399	1.184	-0.013	0.030	0.031	0.044
1	0.683	1.867	-0.010	0.046	0.048	0.058
1	0.821	3.051	-0.012	0.069	0.070	0.082
1	0.889	4.735	-0.019	0.094	0.097	0.116
1	0.925	6.919	-0.031	0.128	0.129	0.160
1	0.947	9.602	-0.048	0.168	0.166	0.214

Here ρ is the unobserved correlation between $Y(1)$ and $Y(0)$; and Δ is the final term in the sample variance equation that we cannot estimate.

Illustration of Neyman Conservative Estimator (Code)

```
n = 100; Y0 = rnorm(n); Y0 = (Y0-mean(Y0))/sd(Y0);  Y1 = rnorm(n) ; Y1=Y1/sd(Y1) + 1
gY1 = function(Y0, Y1,s) {.Y1=Y1+s*Y0; 1+.Y1-mean(.Y1)}
s=(-6:6)/2  #Add + s*Y0 to Y1 for Y1/Y0 covariance
tau = sapply(s, function(i) mean(gY1(Y0, Y1,i) - Y0))
cors= sapply(s, function(i) cor(gY1(Y0, Y1,i),  Y0))  # Note that potential outcomes can be correlated
var1= sapply(s, function(i) var(gY1(Y0, Y1,i))  # Note that potential outcomes can be correlated
phis= sapply(s, function(i) (1/(n-1))*(2*cov(gY1(Y0, Y1,i), Y0) - var(gY1(Y0, Y1,i))-var(Y0)))  # Unknown
tauhat = function(sims, s=1){ .Y1=gY1(Y0, Y1,s)
m = matrix(NA, sims)
for(i in 1:sims){Z=(1:100 %in% sample(n, n/2))
                    m[i]=mean(.Y1[Z] - Y0[!Z])}
m}
neyman = function(sims, s=1, cons=1){.Y1=gY1(Y0, Y1,s)  # Neyman Estimate
 Neyman = matrix(NA, sims)
 for(i in 1:sims){Z=(1:100 %in% sample(n, n/2))
                  Neyman[i]=(n/(n-1))*(var(.Y1[Z])/(n/2) + var(Y0[!Z])/(n/2))+
  (1-cons)*(1/(n-1))*(2*cov(.Y1, Y0) - var(.Y1)-var(Y0))}
mean(Neyman)}
V = sapply(s, function(i) var(tauhat(5000, i)))      # True variance; Empirical estimate
VN1 = sapply(s, function(i) neyman(5000, i, cons=0)) # True variance; Formula check
VN2 = sapply(s, function(i) neyman(5000, i, cons=1)) # Neyman conservative estimate
plot(V, VN1, xlim=c(0, max(VN1, VN2,V)), ylim=c(min(VN2), max(V, VN1, VN2)),
  main="Neyman estimator for  Y(1) and Y(0) correlations, ATE=1, Var(Y(0)=1, Var(Y(1) free" ,
  ylab="Conservative Neyman Estimator", xlab="True Variance", col="grey")
lines(V, VN2, col="red"); text(V, VN2, round(cors,2), offset=TRUE); abline(0,1)
round(cbind(s, tau, cors, var1, phis, V, VN1, VN2), digits=3)
```

Covariate Adjustment

- Even though randomization ensures no bias you may sometimes want to "control" for covariates in order to improve efficiency (see the discussion of blocking above).
- Or you may want to take account of the fact that the assignment to treatment is correlated with a covariate.
- Consider for example this data.
  - You randomly assign offerers to partners in a dictator game (in which offers decide how much of $1 to give to receivers)
  - Your population comes from two groups (80% Baganda and 20% Banyankole) *so in randomly assigning partners you are randomly determining whether a partner is a coethnic or not*
  - **You find that in non coethnic pairings 35% is offered, in coethnic pairings 48% is offered**

Covariate Adjustment

- Your population comes from two groups (80% Baganda and 20% Banyankole) so in randomly assigning partners you are randomly determining whether a partner is a coethnic or not
- You find that in non coethnic pairings 35% is offered, in coethnic pairings 48% is offered
- But a closer look at the data reveals . . .

		To: Baganda	To: Banyankole
Offers by	Baganda	64%	16%
	Banyankole	16%	4%

Number of Games

		To: Baganda	To: Banyankole
Offers by	Baganda	50	50
	Banyankole	20	20

Average Offers

Covariate Adjustment

Control?

- With such data you might be tempted to 'control' for the covariate (here: ethnic group), using regression
- But, perhaps surprisingly, it turns out that regression with covariates does not estimate average treatment effects.
- It does estimate an average of treatment effects, but specifically a minimum variance estimator, not necessarily an estimator of your estimand

Compare:

- $\widehat{\tau}_{ATE} = \sum_x \frac{w_x}{\sum_j w_j} \widehat{\tau}_x$
- $\widehat{\tau}_{OLS} = \sum_x \frac{w_x p_x (1-p_x)}{\sum_j w_j p_j (1-p_j)} \widehat{\tau}_x$

Instead the formula above for $\widehat{\tau}_{ATE}$ is all you need to estimate ATE — at least for discrete covariates.

Covariate Adjustment: Comparison of approaches

Data in Stata

```
* DATA ****************************
set obs 8
g X = _n>4
g T = (_n==1) | (_n>6)
g Y0 = 0
g Y1 = 14*X
g Y = T*Y1 + (1-T)*Y0
g tau = Y1-Y0
egen p = mean(T), by(X)
g ipw = 1/(p*T + (1-p)*(1-T))
list X  Y0 Y1  T Y tau p ipw, mean
```

	X	Y0	Y1	T	Y	tau	p	ipw
1.	0	0	0	1	0	0	.25	4
2.	0	0	0	0	0	0	.25	1.33
3.	0	0	0	0	0	0	.25	1.33
4.	0	0	0	0	0	0	.25	1.33
5.	1	0	14	0	0	14	.5	2
6.	1	0	14	0	0	14	.5	2
7.	1	0	14	1	14	14	.5	2
8.	1	0	14	1	14	14	.5	2
Mean	.5	0	7	.375	3.5	7	.375	2

Randomization Inference

Estimates (see Stata Code)

TRUE ATE	$= 7$
Naive OLS estimate (difference in means)	$= 9.3$ (p $= 0.034$)
OLS estimate controlling for stratum (X)	$= 8.0$ (p $= 0.049$)
Matching estimate (assuming homoskedasticity):	$= 7.0$ (p $= 0.064$)
Regression with inverse propensity score weighting:	$= 7.0$ (p $= 0.207$)
Regression with inverse propensity score weighting & controls:	$= 7.0$ (p $= 0.093$)
t-test of residuals from (Y on $X\|T$) and (Y on $X\|C$):	$= 7.0$ (p $= 0.033$)

One could also run a saturated regression

Covariate Adjustment: Comparison of approaches

Data in Stata

```
* THE TRUTH ***********************************************************
 mean tau

* BIASED APPROACHES: Naive OLS  ***************************************
 reg Y T

* BIASED APPROACHES: OLS Controlling for blocks  **********************
 reg Y T X

* UNBIASED APPROACHES: 1 Matching on stratum (need: ssc install nnmatch)
 nnmatch Y T X

* UNBIASED APPROACHES: 2 Inverse Propensity Score Weighting (IPW)******
 reg Y T [pw = ipw]

* UNBIASED APPROACHES: 3 IPW with controls ****************************
 reg Y T X [pw = ipw]

* UNBIASED APPROACHES: 4 Residual Approach ****************************
 quietly: reg Y X if T==1
 predict R1, res
 quietly: reg Y X if T==0
 predict R0, res
 ttest R1==R0
```

Randomization Inference

- Introducing an entirely new way to think about statistical significance . . .
- Say you randomized assignment to treatment and your data looked like this.

Unit	1	2	3	4	5	6	7	8	9	10
Treatment	0	0	0	0	0	0	0	1	0	0
Healthy?	3	2	4	6	7	2	4	9	8	2

- Does the treatment improve your health?
- $p =?$

Randomization Inference

- Introducing an entirely new way to think about statistical significance . . .
- Say you randomized assignment to treatment and your data looked like this.

Unit	1	2	3	4	5	6	7	8	9	10
Treatment	0	0	0	0	0	0	0	1	0	0
Healthy?	3	2	4	6	7	2	4	8	9	2

- Does the treatment improve your health?
- $p =?$

Randomization Inference

- Say you had a silly randomization procedure and forgot to take account of it in your estimates.

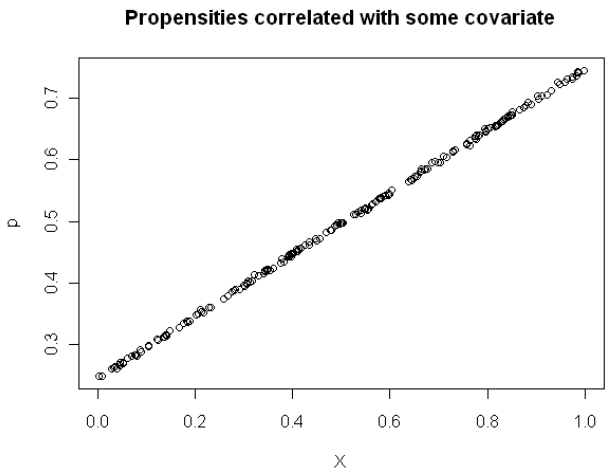

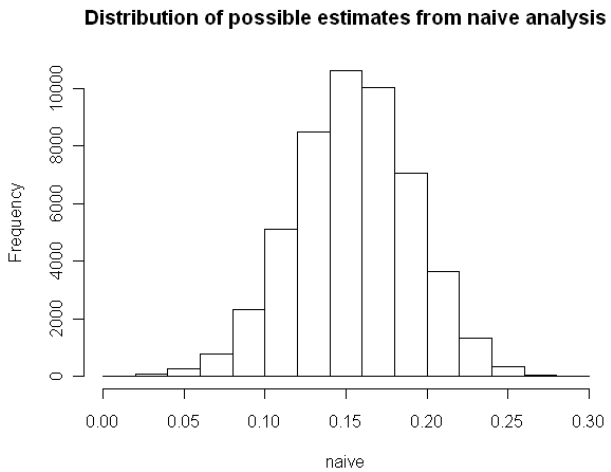

- You estimate .15. *Does the treatment improve your health?*
- $p =?$

LATE—Local Average Treatment Effects

Sometimes you give a medicine but only a non random sample of people actually try to use it. Can you still estimate the medicine's effect?

	$X=0$	$X=1$
$T=0$	$\overline{y}_{00}$ (n_{00})	$\overline{y}_{01}$ (n_{01})
$T=1$	$\overline{y}_{10}$ (n_{10})	$\overline{y}_{11}$ (n_{11})

Say that people are one of 3 types:

- n_a "always takers" have $X=1$ no matter what and have average outcome $\overline{y}_a$
- n_n never takers have $X=0$ no matter what with outcome $\overline{y}_n$
- n_c compliers have $X=T$ and average outcomes $\overline{y}_c^1$ if treated and $\overline{y}_c^0$ if not.

	$X=0$	$X=1$
$T=0$	$\frac{\frac{1}{2}n_c}{\frac{1}{2}n_c+n_{10}}\overline{y}_c^0+\frac{n_{10}}{\frac{1}{2}n_c+n_{10}}\overline{y}_{10}$	$\overline{y}_{01}=\overline{y}_a$
$T=1$	$\overline{y}_{10}=\overline{y}_n$	$\frac{\frac{1}{2}n_c}{\frac{1}{2}n_c+n_{01}}\overline{y}_c^1+\frac{n_{01}}{\frac{1}{2}n_c+n_{01}}\overline{y}_{01}$

LATE—Local Average Treatment Effects

You give a medicine to 50% but only a non random sample of people actually try to use it. Can you still estimate the medicine's effect?

	$X=0$	$X=1$
$T=0$	$\frac{\frac{1}{2}n_c}{\frac{1}{2}n_c+n_{10}}\overline{y}^0_c + \frac{n_{10}}{\frac{1}{2}n_c+n_{10}}\overline{y}_{10}$	$\overline{y}_{01} = \overline{y}_a$
$T=1$	$\overline{y}_{10} = \overline{y}_n$	$\frac{\frac{1}{2}n_c}{\frac{1}{2}n_c+n_{01}}\overline{y}^1_c + \frac{n_{01}}{\frac{1}{2}n_c+n_{01}}\overline{y}_{01}$

$$\overline{y}^1_c - \overline{y}^0_c = (\overline{y}_{11} - \overline{y}_{00}) + 2\frac{n_{01}}{n_c}(\overline{y}_{11} - \overline{y}_{01}) + 2\frac{n_{10}}{n_c}(\overline{y}_{10} - \overline{y}_{00})$$

$$\text{Average in } T=1 \text{ group} = \frac{\left(n_{10}\overline{y}_{10} + \frac{1}{2}n_c\overline{y}^1_c\right) + n_{01}\overline{y}_{01}}{n_{01}+n_{10}+\frac{1}{2}n_c}$$

$$\text{Average in } T=0 \text{ group} = \frac{n_{10}\overline{y}_{10} + \left(\frac{1}{2}n_c\overline{y}^0_c + n_{01}\overline{y}_{01}\right)}{n_{01}+n_{10}+\frac{1}{2}n_c}$$

$$\text{Difference} = (\overline{y}^1_c - \overline{y}^0_c)\frac{n_c}{n} \text{ So: } LATE = ITT \times \frac{n}{n_c}$$

SUTVA violations (Spillovers)

Spillovers can result in the estimation of weaker effects in cases where effects are actually stronger.

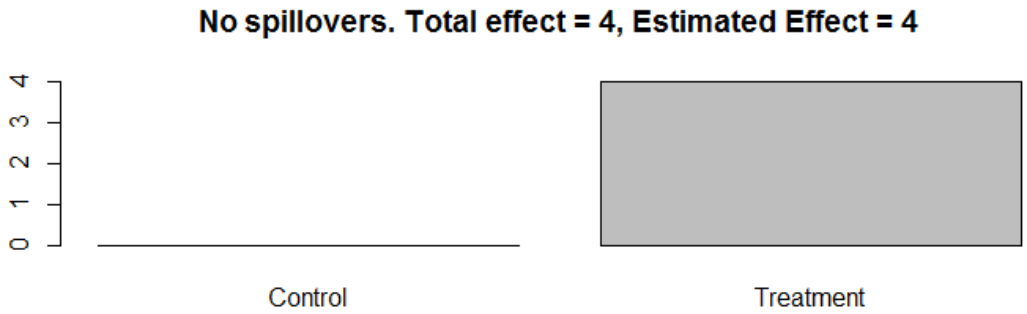

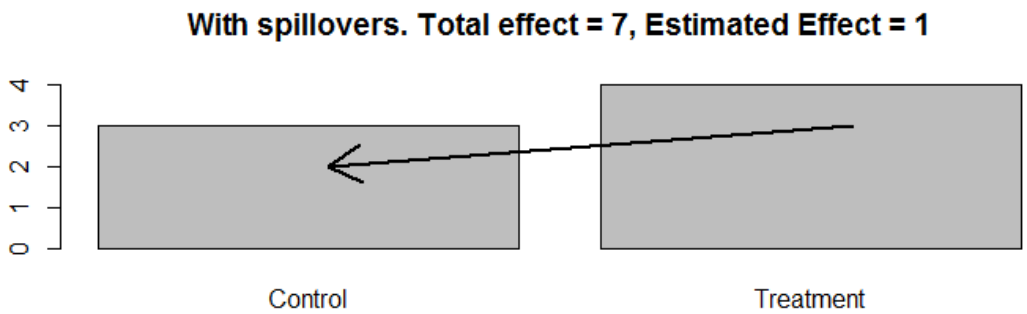

SUTVA violations (Spillovers)

```
par(mfrow=c(2,1)); X=c("Control", "Treatment")
barplot(c(0,4), names.arg=X, main="No spillovers. Total effect = 4, Estimated Effect = 4");
barplot(c(3,4), names.arg=X, main="With spillovers. Total effect = 7, Estimated Effect = 1")
arrows(2,3,.75,2, lwd=2)
```

SUTVA violations

- In the presence of spillovers outcomes depend not just on own assignment but also on the assignment of others.
- May have different potential outcomes for every possible profile

Unit	Location	$D_\emptyset$	$y(D_\emptyset)$	D_1	$y(D_1)$	D_2	$y(D_2)$	D_3	$y(D_3)$	D_4	$y(D_4)$
A	1	0	0	1	3	0	1	0	0	0	0
B	2	0	0	0	3	1	3	0	3	0	0
C	3	0	0	0	0	0	3	1	3	0	3
D	4	0	0	0	0	0	0	0	1	1	3
$\bar{y}_{\text{treated}}$			-		3		3		3		3
$\bar{y}_{\text{untreated}}$			0		1		4/3		4/3		1
$\bar{y}_{\text{neighbors}}$			-		3		2		2		3
$\bar{y}_{\text{pure control}}$			0		0		0		0		0
ATT (direct effect)			-		3		3		3		3
ATT (indirect effect)			-		3		2		2		3

Potential outcomes for four units for different treatment profiles, D_1-D_4. D_i represents an allocation to treatment and $y_j(D_i)$ is the potential outcome for (row) unit j given (column) allocation i.
Assumption: Spillovers only affect immediate neighbors.

SUTVA violations

Unit	Location	$D_\emptyset$	$y(D_\emptyset)$	D_1	$y(D_1)$	D_2	$y(D_2)$	D_3	$y(D_3)$	D_4	$y(D_4)$
A	1	0	0	1	3	0	1	0	0	0	0
B	2	0	0	0	3	1	3	0	3	0	0
C	3	0	0	0	0	0	3	1	3	0	3
D	4	0	0	0	0	0	0	0	1	1	3

Potential outcomes for four units for different treatment profiles, D_1-D_4. D_i represents an allocation to treatment and $y_j(D_i)$ is the potential outcome for (row) unit j given (column) allocation i.

- The key is to think through the structure of spillovers.
- Here immediate neighbors are exposed
- In this case we can define a direct treatment (being exposed) and an indirect treatment (having a neighbor exposed) and we can work out *the propensity for each unit of receiving each type of treatment*
- These may be non uniform (here central types are more likely to have teated neighbors); but we can still use the randomization to assess effects

Hard Limits

- Real Time
- History has Happened
- Power & Scale
- Variables as Attributes
- The assignment process matters
- Chronic spillovers
- External validity
- Ethics