Lab 3

Olivier Bergeron-Boutin

February 2nd, 2021

Introductory meme

Colchicine

« On vient de fournir à la planète un espoir ! s'exclame au bout du fil le Dr Jean-Claude Tardif, chercheur principal de l'étude COLCORONA et directeur du centre de recherche de l'Institut de cardiologie de Montréal (ICM). On a finalement un premier traitement qui peut aider les patients atteints de la COVID-19 avant leur admission à l'hôpital pour prévenir les hospitalisations, prévenir les intubations et prévenir les décès. »

Chez 4159 patients qui présentaient un facteur de risque de complications et dont le diagnostic de COVID-19 avait été validé par un test PCR, la colchicine a entraîné une baisse des hospitalisations de 25 %, une baisse du besoin de ventilation de 50 % et une diminution des décès de 44 % par rapport au groupe témoin. « C'est une percée majeure », déclare le Dr Tardif.

Colchicine

« On vient de fournir à la planète un espoir ! s'exclame au bout du fil le D^r^ Jean-Claude Tardif, chercheur principal de l'étude COLCORONA et directeur du centre de recherche de l'Institut de cardiologie de Montréal (ICM). On a finalement un premier traitement qui peut aider les patients atteints de la COVID-19 avant leur admission à l'hôpital pour prévenir les hospitalisations, prévenir les intubations et prévenir les décès. »

Chez 4159 patients qui présentaient un facteur de risque de complications et dont le diagnostic de COVID-19 avait été validé par un test PCR, la colchicine a entraîné une baisse des hospitalisations de 25 %, une baisse du besoin de ventilation de 50 % et une diminution des décès de 44 % par rapport au groupe témoin. « C'est une percée majeure », déclare le D^r^ Tardif.

Hospitalizations: $\downarrow$25%

Mechanical ventilation: $\downarrow$50%

Deaths: $\downarrow$44%

Colchicine

« On vient de fournir à la planète un espoir ! s'exclame au bout du fil le Dr Jean-Claude Tardif, chercheur principal de l'étude COLCORONA et directeur du centre de recherche de l'Institut de cardiologie de Montréal (ICM). On a finalement un premier traitement qui peut aider les patients atteints de la COVID-19 avant leur admission à l'hôpital pour prévenir les hospitalisations, prévenir les intubations et prévenir les décès. »

Chez 4159 patients qui présentaient un facteur de risque de complications et dont le diagnostic de COVID-19 avait été validé par un test PCR, la colchicine a entraîné une baisse des hospitalisations de 25 %, une baisse du besoin de ventilation de 50 % et une diminution des décès de 44 % par rapport au groupe témoin. « C'est une percée majeure », déclare le Dr Tardif.

Hospitalizations: ↓25%

Mechanical ventilation: ↓50%

Deaths: ↓44%

But...no data, no pre-print - just a press release that says the results "approached statistical significance"

Colchicine

Table 2. Rates and Odds Ratios for Major Clinical Outcomes.

Clinical Outcome	Colchicine	Placebo	Odds Ratio (95% CI)	P Value
ITT population	N=2235	N=2253		
Primary composite endpoint - no. (%)	104 (4.7%)	131 (5.8%)	0.79 (0.61-1.03)	0.08
Components of primary endpoint:				
Death - no. (%)	5 (0.2%)	9 (0.4%)	0.56 (0.19-1.67)	
Hospitalization for COVID-19 no. (%)	101 (4.5%)	128 (5.7%)	0.79 (0.60-1.03)	
Secondary endpoint:				
Mechanical ventilation - no. (%)	11 (0.5%)	21 (0.9%)	0.53 (0.25-1.09)	
Patients with PCR-proven COVID-19	N=2075	N=2084		
Primary composite endpoint – no. (%)	96 (4.6%)	126 (6.0%)	0.75 (0.57-0.99)	0.04
Components of primary endpoint:				
Death – no. (%)	5 (0.2%)	9 (0.4%)	0.56 (0.19-1.66)	
Hospitalization for COVID-19 no. (%)	93 (4.5%)	123 (5.9%)	0.75 (0.57-0.99)	
Secondary endpoint:				
Mechanical ventilation – no. (%)	10 (0.5%)	20 (1.0%)	0.50 (0.23-1.07)	

Colchicine

Table 1. Characteristics of the Trial Patients.

Characteristic	Colchicine (N=2235)	Placebo (N=2253)
Age - years	54.4±9.7	54.9±9.9
Female sex - no. (%)	1238 (55.4%)	1183 (52.5%)
Caucasian - no. (%)	2086 (93.3%)	2096 (93.2%)
Body-mass index (kg/m^2)	30.0±6.2	30.0±6.3
Smoking - no. (%)	217 (9.7%)	212 (9.4%)

"Our trial has certain limitations. The study was stopped when 75% of the planned patients were recruited and had completed the 30-day follow-up. In addition to the logistical issues faced in the current challenging context, the perceived need to disseminate the study results rapidly in view of the current state of the pandemic largely contributed to our decision. The duration of follow-up was relatively short at approximately 30 days."

Colchicine

“Our trial has certain limitations. The study was stopped when 75% of the planned patients were recruited and had completed the 30-day follow-up. In addition to the logistical issues faced in the current challenging context, the perceived need to disseminate the study results rapidly in view of the current state of the pandemic largely contributed to our decision. The duration of follow-up was relatively short at approximately 30 days.”

Did they know the interim results when the study was interrupted?


PRESS RELEASE
ABSTRACT
RESULTS

Block randomization

Wantchekon (2003) convinced the campaigns of major presidential candidates in Benin to randomize the messages they employed in 24 villages. We're are interested in two treatments:

Block randomization

Wantchekon (2003) convinced the campaigns of major presidential candidates in Benin to randomize the messages they employed in 24 villages. We're are interested in two treatments:

- Public policy message: emphasized "national unity and peace, eradicating corruption, alleviating poverty..."

Block randomization

Wantchekon (2003) convinced the campaigns of major presidential candidates in Benin to randomize the messages they employed in 24 villages. We're are interested in two treatments:

- Public policy message: emphasized "national unity and peace, eradicating corruption, alleviating poverty..."
- Clientelist message: "a specific promise to the village" such as "government patronage jobs or local public goods"

Block randomization

Wantchekon (2003) convinced the campaigns of major presidential candidates in Benin to randomize the messages they employed in 24 villages. We're are interested in two treatments:

- Public policy message: emphasized "national unity and peace, eradicating corruption, alleviating poverty..."
- Clientelist message: "a specific promise to the village" such as "government patronage jobs or local public goods"

(Think about how insane this is!)

Block randomization

Wantchekon (2003) convinced the campaigns of major presidential candidates in Benin to randomize the messages they employed in 24 villages. We're are interested in two treatments:

- Public policy message: emphasized "national unity and peace, eradicating corruption, alleviating poverty..."
- Clientelist message: "a specific promise to the village" such as "government patronage jobs or local public goods"

(Think about how insane this is!)

Our data is at the village-level

Block randomization

Wantchekon (2003) convinced the campaigns of major presidential candidates in Benin to randomize the messages they employed in 24 villages. We're are interested in two treatments:

- Public policy message: emphasized "national unity and peace, eradicating corruption, alleviating poverty..."
- Clientelist message: "a specific promise to the village" such as "government patronage jobs or local public goods"

(Think about how insane this is!)

Our data is at the village-level

Block randomization: 8 blocks of 3 villages based on geography

Block randomization

Wantchekon (2003) convinced the campaigns of major presidential candidates in Benin to randomize the messages they employed in 24 villages. We're are interested in two treatments:

- Public policy message: emphasized "national unity and peace, eradicating corruption, alleviating poverty..."
- Clientelist message: "a specific promise to the village" such as "government patronage jobs or local public goods"

(Think about how insane this is!)

Our data is at the village-level

Block randomization: 8 blocks of 3 villages based on geography

Our outcome is `vote_pop` and our treatment `treat`

Wantchekon: 1a

Estimate the ATE (point estimate, standard error and p-value) of the clientelist message compared to the public policy message, using a regression or difference in means estimator, but ignoring the blocking.

Wantchekon: 1a

Estimate the ATE (point estimate, standard error and p-value) of the clientelist message compared to the public policy message, using a regression or difference in means estimator, but ignoring the blocking.

```
treat <- filter(want, treat == "client")
control <- filter(want, treat == "pub.pol")
ate <- mean(treat$vote_pop - mean(control$vote_pop))
ate
```

```
## [1] 0.1575
```

```
se <- sqrt(var(treat$vote_pop)/8 + var(control$vote_pop)/8)
se
```

```
## [1] 0.08885934
```

ATE with regression

```r
library(sandwich)
library(lmtest)
m1 <- lm(vote_pop ~ treat, want)
coeftest(m1, vcov = vcovHC(m1, type = "HC1"))
```

```
##
## t test of coefficients:
##
##                Estimate Std. Error t value  Pr(>|t|)
## (Intercept)    0.843750   0.032068 26.3114 2.541e-13 ***
## treatpub.pol -0.157500   0.088859 -1.7725   0.09807 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Wantchekon: 1b

Estimate the ATE of the clientelist message compared to the public policy message, using a regression estimator that takes the blocking into account. How do the results differ and why?

```
m2 <- lm(vote_pop ~ treat + factor(block), want)
coeftest(m2, vcov = vcovHC(m2, type = "HC1"))
```

```
##
## t test of coefficients:
##
##                  Estimate Std. Error t value  Pr(>|t|)
## (Intercept)      0.783750   0.042276 18.5388 3.296e-07 ***
## treatpub.pol    -0.157500   0.063238 -2.4906  0.041562 *
## factor(block)2   0.170000   0.063948  2.6584  0.032541 *
## factor(block)3   0.230000   0.083687  2.7483  0.028573 *
## factor(block)4  -0.175000   0.216968 -0.8066  0.446447
## factor(block)5   0.200000   0.083687  2.3898  0.048180 *
## factor(block)6   0.085000   0.029580  2.8735  0.023871 *
## factor(block)7   0.115000   0.041619  2.7632  0.027968 *
## factor(block)8  -0.145000   0.030531 -4.7493  0.002085 **
## ---
```

Wantchekon: 1c

In this question, you will test the sharp null that there is no treatment effect using randomization inference (Fisher's Exact Test) and ignore the blocking. In other words, pretend that the treatment was assigned using complete randomization. Use the difference in means as your test-statistic. (If you don't want to enumerate every possible treatment assignment, just sample a large number of draws from the randomization distribution.) Please answer the following:

In this question, you will test the sharp null that there is no treatment effect using randomization inference (Fisher's Exact Test) and ignore the blocking. In other words, pretend that the treatment was assigned using complete randomization. Use the difference in means as your test-statistic. (If you don't want to enumerate every possible treatment assignment, just sample a large number of draws from the randomization distribution.) Please answer the following:
How do you interpret the null hypothesis for this test?

Wantchekon: 1c

In this question, you will test the sharp null that there is no treatment effect using randomization inference (Fisher's Exact Test) and ignore the blocking. In other words, pretend that the treatment was assigned using complete randomization. Use the difference in means as your test-statistic. (If you don't want to enumerate every possible treatment assignment, just sample a large number of draws from the randomization distribution.) Please answer the following:
How do you interpret the null hypothesis for this test?

$Y_i(1) = Y_i(0) \forall i$

Wantchekon: 1c

In this question, you will test the sharp null that there is no treatment effect using randomization inference (Fisher's Exact Test) and ignore the blocking. In other words, pretend that the treatment was assigned using complete randomization. Use the difference in means as your test-statistic. (If you don't want to enumerate every possible treatment assignment, just sample a large number of draws from the randomization distribution.) Please answer the following:
How do you interpret the null hypothesis for this test?

$$Y_i(1) = Y_i(0) \forall i$$
$$Y_i(1) - Y_i(0) = 0$$

Wantchekon: 1c

In this question, you will test the sharp null that there is no treatment effect using randomization inference (Fisher's Exact Test) and ignore the blocking. In other words, pretend that the treatment was assigned using complete randomization. Use the difference in means as your test-statistic. (If you don't want to enumerate every possible treatment assignment, just sample a large number of draws from the randomization distribution.) Please answer the following:
How do you interpret the null hypothesis for this test?

$$Y_i(1) = Y_i(0) \forall i$$
$$Y_i(1) - Y_i(0) = 0$$

Standard null: merely about the *expectation*: $\mathbb{E}[Y_i(1)] = \mathbb{E}[Y_i(0)]$

Wantchekon: 1d

Plot the distribution of the treatment effect under the sharp null, and show where on this distribution the realized test-statistic falls. What is the p-value you would obtain, and how does it differ from the one you obtained in 1a. Which do you prefer?

Wantchekon: 1d

Plot the distribution of the treatment effect under the sharp null, and show where on this distribution the realized test-statistic falls. What is the p-value you would obtain, and how does it differ from the one you obtained in 1a. Which do you prefer?

```
avetreat <- c()
set.seed(0202)
for(i in 1:1000){
  s <- sample(16, replace = FALSE)
  control <- want[s[1:8],]
  treat <- want[s[9:16],]
  avetreat[i] <- mean(treat$vote_pop - mean(control$vote_pop))
}
sum(abs(avetreat) >= abs(ate)) / length(avetreat)
```

```
## [1] 0.091
```

1d (cont)

```r
hist(avetreat)
abline(v = ate, col = "red")
```

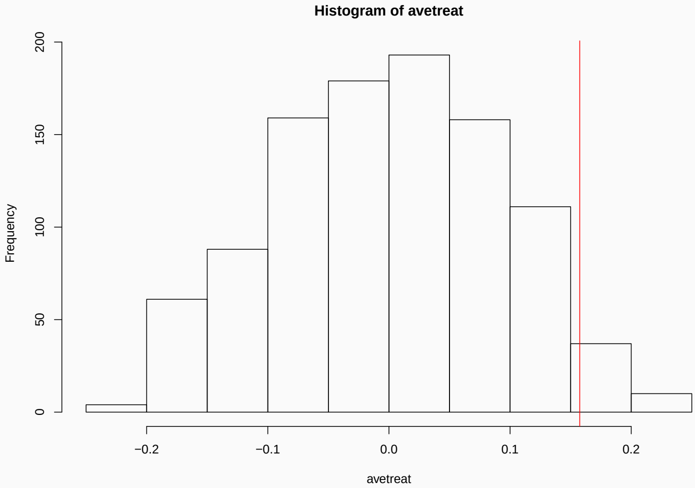

Repeat 1d, but taking blocking into account

Wantchekon: 1e

Repeat 1d, but taking blocking into account

```
avetreat_block <- c()
set.seed(0202)
for(i in 1:1000){
  want$treatind <- 0
  for(j in 1:8){
    ran <- runif(1)
    if(ran >= 0.5){
      want[2*j,6] <- 1
    }else{
      want[(2*j-1),6] <- 1
    }
  }
  control <- filter(want, treatind == 0)
  treat <- filter(want, treatind == 1)
  avetreat_block[i] <- mean(treat$vote_pop - mean(control$vote_pop))
}
sum(abs(avetreat_block) >= abs(ate)) / length(avetreat_block)
```

```
## [1] 0.005
```

1f

What can you conclude about the effectiveness of clientelistic appeals in Benin? For extra credit, comment on the internal and external validity of Wantchekon's experiment (you may need to read the full article for this). Was his choice of units of analysis and units of randomization justified?

Cluster randomization