A **CACE** in point: Estimating causal effects via a latent class approach in RCTs with noncompliance using Stata

Stata Conference 2021

Dr Patricio Troncoso
Heriot-Watt University
University of Manchester (Honorary)
@p_tronc
p.troncoso@hw.ac.uk

2021/08/05

Outline

- RCTs under noncompliance and CACE.
- Using Stata **gsem** to estimate a CACE model.
- Comparison with other software packages.
- A recent application: The Good Behaviour Game (school-based RCT) in England.

Randomised controlled trials (RCT) under noncompliance

A standard RCT compares the outcome of interest in two groups:

- The treatment group, who is randomly selected to receive the treatment
- The control group, who is randomly selected to continue with standard treatment or no treatment at all

In the presence of perfect compliance, a dummy for trial arm assignment is enough to estimate the effect of the intervention.

- That is known as "Intention to Treat" (ITT), where we use standard regression models
- But some participants may not receive the treatment in full or comply with all requirements
- As expected, this is a common occurrence that can bias the ITT estimate

What is CACE?

Complier Average Causal Effect -> Proposed by Imbens and Rubin (1997)

The difference between the outcome in those participants who complied with the intervention and those who would have complied if assigned to treatment

	Compliance = 1	Compliance = 0
Treatment = 1	a) Complier	b) Non-complier
Treatment = 0	c) "Would-be" complier	d) Non-complier

What is CACE?

Complier Average Causal Effect -> Proposed by Imbens and Rubin (1997)

The difference between the outcome in those participants who complied with the intervention and those who would have complied if assigned to treatment

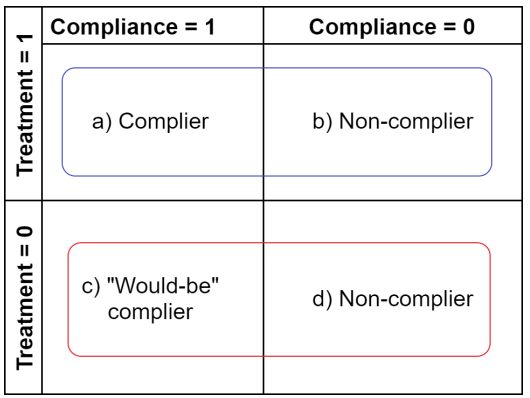

- ITT = (a + b) - (c + d)

What is CACE?

Complier Average Causal Effect -> Proposed by Imbens and Rubin (1997)

The difference between the outcome in those participants who complied with the intervention and those who would have complied if assigned to treatment

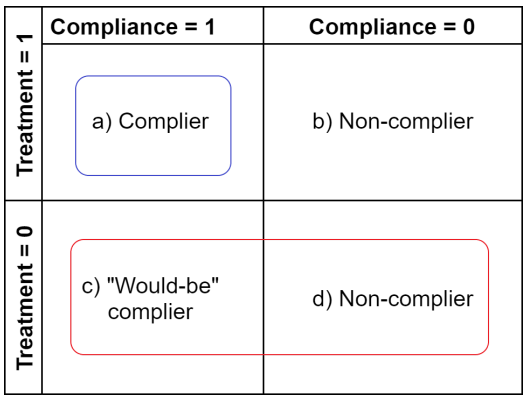

- ITT = (a + b) - (c + d)
- Per protocol = a - (c + d)

What is CACE?

Complier Average Causal Effect -> Proposed by Imbens and Rubin (1997)

The difference between the outcome in those participants who complied with the intervention and those who would have complied if assigned to treatment

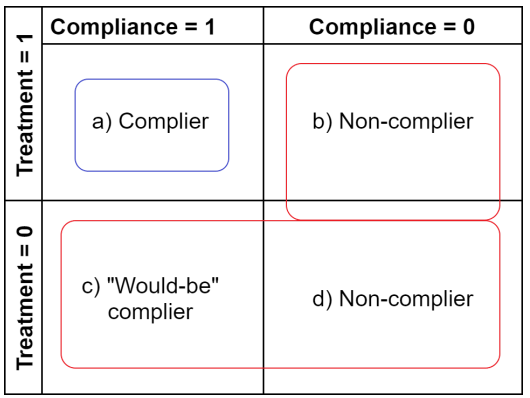

- ITT = (a + b) - (c + d)
- Per protocol = a - (c + d)
- As treated = a - (b + c + d)

What is CACE?

Complier Average Causal Effect -> Proposed by Imbens and Rubin (1997)

The difference between the outcome in those participants who complied with the intervention and those who would have complied if assigned to treatment

	Compliance = 1	Compliance = 0
Treatment = 1	a) Complier	b) Non-complier
Treatment = 0	c) "Would-be" complier	d) Non-complier

- ITT = (a + b) - (c + d)
- Per protocol = a - (c + d)
- As treated = a - (b + c + d)
- CACE = a - c

What is CACE?

Complier Average Causal Effect -> Proposed by Imbens and Rubin (1997)

The difference between the outcome in those participants who complied with the intervention and those who would have complied if assigned to treatment

	Compliance = 1	Compliance = 0
Treatment = 1	a) Complier	b) Non-complier
Treatment = 0	c) "Would-be" complier	d) Non-complier

- ITT = (a + b) - (c + d)
- Per protocol = a - (c + d)
- As treated = a - (b + c + d)
- CACE = a - c

But compliance is not observable in the control group, which is why:

- a) we need to make assumptions exploiting randomisation, and
- b) adopt a probabilistic approach

Assumptions of CACE

Random assignment

- This allows us to assume "equivalent" groups with respect to the outcome before the intervention

Monotonicity

- There are no *defiers* (those who do the opposite of assignment)
- Similarly, there are no *always-takers* (those who participate regardless of assignment)

Stable unit treatment value (SUTVA)

- The outcome for any participant is independent of the group assignment of other participants

Exclusion restriction

- The treatment has **zero** effect on the outcome in the non-complier group

What are the characteristics of a CACE statistical model?

To estimate a CACE model, we need the following:

- An appropriate regression model for the outcome of interest
- A binary variable indicating compliance in the treatment arm that is treated as known in the model.
- A probabilistic model for compliance in the control group with reasonable predictors

CACE and latent class models: An example

"JOBS II": RCT to prevent depression as a result of job loss: **Little & Yau (1998)**. Data available **here**

Our outcome of interest is **"depress"** and **"r"** is the treatment dummy.

"c" is a dummy to indicate compliance. From this, we derive **"comp"** as compliance with the intervention in the treatment arm and leave it as missing (unobserved) in the control group.

There are also a few covariates, such as age and education level, etc.

We could fit a standard latent class regression model for *depress* with latent compliance as such:

```
gsem (depress <- r depbase risk) ///
  (C <- age educ motivate econ assert single nonwhite) ///
  (comp <- , logit), ///
  lclass(C 2)
```

But this is not a CACE model because of the following:

- The treatment effect (**"r"**) is estimated freely
- The observed compliance in the treatment arm is not treated as *known*

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (depress <- r depbase risk) ///
  (C <- age educ motivate econ assert single nonwhite) ///
  (comp <- , logit), ///
  lclass(C 2)
```

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (depress <- r depbase risk) ///
  (C <- age educ motivate econ assert single nonwhite) ///
  (comp <- , logit), ///
  lclass(C 2)
```

Step 1: extend the **regression model for depression**

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (1.C: depress <- r depbase risk) ///
  (2.C: depress <- r depbase risk) ///
  (comp <- , logit) ///
  (C <- age educ motivate econ assert single nonwhite), ///
  lclass(C 2)
```

Step 1: extend the regression model for depression **into 2 paths**

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (1.C: depress <- r depbase risk) ///
  (2.C: depress <- r depbase risk) ///
  (comp <- , logit) ///
  (C <- age educ motivate econ assert single nonwhite), ///
  lclass(C 2)
```

Step 1: extend the regression model for depression into 2 paths

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (1.C: depress <- r depbase risk) ///
  (2.C: depress <- r depbase risk) ///
  (comp <- , logit) ///
  (C <- age educ motivate econ assert single nonwhite), ///
  lclass(C 2)
```

Step 1: extend the regression model for depression into 2 paths

Step 2: fix the effect of the intervention in the **non-compliers class**

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (1.C: depress <- r@0 depbase risk) ///
  (2.C: depress <- r depbase risk) ///
  (comp <- , logit) ///
  (C <- age educ motivate econ assert single nonwhite), ///
  lclass(C 2)
```

Step 1: extend the regression model for depression into 2 paths

Step 2: fix the effect of the intervention in the **non-compliers class** to **zero**:

This is the **"exclusion restriction"** assumption

- `1.C`: regression path for non-compliers. Treatment effect fixed at zero (`@0`)
- `2.C`: regression path for compliers. Treatment effect freely estimated.

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (1.C: depress <- r@0 depbase risk) ///
  (2.C: depress <- r depbase risk) ///
  (comp <- , logit) ///
  (C <- age educ motivate econ assert single nonwhite), ///
  lclass(C 2)
```

Step 1: extend the regression model for depression into 2 paths

Step 2: fix the effect of the intervention in the **non-compliers class** to **zero**:

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (1.C: depress <- r@0 depbase risk) ///
  (2.C: depress <- r depbase risk) ///
  (comp <- , logit) ///
  (C <- age educ motivate econ assert single nonwhite), ///
  lclass(C 2)
```

Step 1: extend the regression model for depression into 2 paths

Step 2: fix the effect of the intervention in the **non-compliers class** to **zero**:

Step 3: extend the **latent class model** for compliance

Making the most of gsem path notation

We can use gsem path notation to make a standard LC model into a CACE model:

```
gsem (1.C: depress <- r@0 depbase risk) ///
  (2.C: depress <- r depbase risk) ///
  (1.C: comp <- _cons@-15, logit) ///
  (2.C: comp <- _cons@15, logit) ///
  (C <- age educ motivate econ assert single nonwhite), ///
  lclass(C 2)
```

Step 1: extend the regression model for depression into 2 paths

Step 2: fix the effect of the intervention in the **non-compliers class** to **zero**:

Step 3: extend the **latent class model** for compliance into 2 paths to treat observed compliance in the treatment arm as *known*

- `1.C`: path for non-compliers. Intercept fixed at `logit=-15` (`_cons@-15`) to ensure non-compliers in the treatment arm (`comp=0`) will be assigned to that class
- `2.C`: path for compliers. Intercept fixed at `logit=15` (`_cons@15`) to ensure compliers in the treatment arm (`comp=1`) will be assigned to that class.

Estimation and results: JOBS II dataset

This code is to replicate the results of Little & Yau (1998) using gsem (Stata 15 onwards):

```
/* Read the data */

infile depress risk r depbase age motivate educ assert single econ nonwhite x10 c c0 ///
using "http://www.gllamm.org/books/wjobs.dat", clear

/* Generate compliance indicator variable */

gen comp=c if r==1 /* compliance missing in control group */

/* Fit the model

gsem (1.C: depress <- r@0 depbase@c1 risk@c2) ///
  (2.C: depress <- r depbase@c1 risk@c2) ///
  (1.C: comp <- _cons@-15, logit) ///
  (2.C: comp <- _cons@15, logit) ///
  (C <- age educ motivate econ assert single nonwhite), ///
  lclass(C 2)
```

Note: `@c1` and `@c2` constrain the effect of `depbase` and `risk` on `depress` to equality across classes

Estimation and results: JOBS II dataset

These are the results for the Non-compliers class (class 1)

```
Class          : 1

Response       : depress                        Number of obs     =        502
Family         : Gaussian
Link           : identity

Response       : comp                           Number of obs     =        335
Family         : Bernoulli
Link           : logit
```

	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
depress						
depbase	-1.463379	.1826867	-8.01	0.000	-1.821438	-1.10532
r	0	(omitted)				
risk	.9117568	.2624529	3.47	0.001	.3973586	1.426155
_cons	1.632537	.2791255	5.85	0.000	1.085461	2.179613
comp						
_cons	-15	(constrained)				
var(e.depress)	.506397	.0322776			.4469262	.5737814

The exclusion restriction for non-compliers is applied and the effect of the intervention ("**r**") on the outcome in this group is fixed to **zero**

Estimation and results: JOBS II dataset

These are the results for the Compliers class (class 2)

```
Class          : 2

Response       : depress                          Number of obs     =        502
Family         : Gaussian
Link           : identity

Response       : comp                             Number of obs     =        335
Family         : Bernoulli
Link           : logit
```

	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
depress						
depbase	-1.463379	.1826867	-8.01	0.000	-1.821438	-1.10532
r	-.3098673	.1173219	-2.64	0.008	-.5398141	-.0799205
risk	.9117568	.2624529	3.47	0.001	.3973586	1.426155
_cons	1.81249	.2971227	6.10	0.000	1.23014	2.394839
comp						
_cons	15	(constrained)				
var(e.depress)	.506397	.0322776			.4469262	.5737814

The effect of treatment in the compliers class is **-0.3098673**

- This means that *compliers* are expected to score *0.31* less in the depression scale than *"would-be"* compliers

Other software packages

gllamm (Rabe-Hesketh, S. et al., 2004). JOBS II example available here

Mplus (Muthén, L., & Muthén, B., 1998-2017). JOBS II example available here

Latent Gold (Vermunt J., & Magidson, J., 2016). JOBS II example comes with demo.

gsem

Coef.	Std. Err.	z	P>|z|
-.3098673	.1173219	-2.64	0.008

gllamm

Coef.	Std. Err.	z	P>|z|
-.3098673	.1173219	-2.64	0.008

Mplus

	Estimate	S.E.	Est./S.E.	P-Value
Latent Class 1				
DEPRESS ON				
TX	-0.310	0.130	-2.378	0.017
RISK	0.912	0.247	3.685	0.000
DEPBASE	-1.463	0.181	-8.077	0.000

Latent Gold

term		coef	s.e.	z-value	p-value
treatment	| Class(1)	-0.3099	0.1173	-2.6412	0.0083
treatment	| Class(2)	0	.	.	.
risk		0.9118	0.2625	3.474	0.00051
depbase		-1.4634	0.1827	-8.0103	1.10E-15

A recent application: The Good Behaviour Game

The Good Behaviour Game

GBG is a school-based, universal behaviour management intervention implemented by class teachers.

- Core components: classroom rules, team membership, monitoring behaviour and positive reinforcement.
- Children work in teams to win the GBG for agreed rewards.
- Played during normal classroom activities for a specified time period.
- The class teacher monitors rules: working quietly, being polite, etc.

GBG trial in England

- 3084 children in 77 schools
- 1524 children in 39 control schools
- 1560 children in 38 intervention schools
- Intervention: 2015-2017
- Follow-up: 2017-2019
- Primary data collection
- Administrative data: National Pupil Database

Pre-registered hypotheses

Children in English primary schools implementing the GBG would demonstrate significantly **better**:

- mental health (conduct problems, psychological wellbeing, etc.);
- rates of **school absence** from school (the focus of this presentation)

For more details:

The GBG trial registration is available here

The full study protocol is available here

The final report is currently under review and will be published in the Journal of Public Health Research

Compliance with the GBG protocol

The duration of GBG play sessions varied across classrooms widely, so we chose *dosage* as our main indicator of compliance:

- **"Compliers"** are those children in classrooms where the GBG was played for **at least the median time** across intervention schools
- **"Non-compliers"** are those in classrooms where the GBG was played for **less than the median time**
- We also used the **75th percentile** as sensitivity analysis

CACE model for **absence from school** in the GBG trial (1)

We used Stata gsem to run a CACE model for absence

```
Class          : 1

Response       : absence17                    Number of obs     =      2,888
Family         : Poisson
Link           : log

Response       : c50                          Number of obs     =      1,308
Family         : Bernoulli
Link           : logit

                                  (Std. Err. adjusted for 76 clusters in proj_schid)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
absence17    |
    baseline |   13.29833   1.31642    10.10   0.000     10.7182    15.87847
      gender |  -.0150795   .0643436    -0.23   0.815    -.1411906    .1110317
         fsm |   .2514962   .1016229     2.47   0.013     .052319    .4506735
         sen |    .172562   .1119566     1.54   0.123    -.0468688    .3919929
         ks1 |  -.0320073   .0161512    -1.98   0.048    -.0636631   -.0003516
      trial2 |          0  (omitted)
  condpr1_tot |  -.0141845   .0161751    -0.88   0.381    -.0458871    .0175181
concentration1|  -.0235934   .0361202    -0.65   0.514    -.0943876    .0472009
   prosocial1 |   .031603    .04673      0.68   0.499     -.059986    .123192
  school_size |   -.00003   .0002625    -0.11   0.909    -.0005446    .0004845
   school_fsm |  -.0017833   .0053342    -0.33   0.738    -.0122383    .0086716
  school_lev4 |   .005011    .004732     1.06   0.290    -.0042636    .0142855
       schcond|   .1048374   .102594     1.02   0.307    -.0962431    .3059179
   school_eal |   .0030782   .0030872     1.00   0.319    -.0029726     .009129
        _cons |  -4.460392   .6226664    -7.16   0.000    -5.680795   -3.239988
ln(sessions17)|          1  (exposure)
-------------+----------------------------------------------------------------
c50          |
        _cons |        -15  (constrained)
------------------------------------------------------------------------------
```

- This is the non-compliers class. The effect of **trial** is fixed to *zero*.
- **Absence**. Poisson distribution: `family(poisson) exposure(exposure_var)`
- **Clustering around schools**. Robust standard errors: `vce(cluster clust_var)`

CACE model for **absence from school** in the GBG trial (2)

The results for the compliers class are:

```
Class          : 2

Response       : absence17                    Number of obs     =      2,888
Family         : Poisson
Link           : log

Response       : c50                          Number of obs     =      1,308
Family         : Bernoulli
Link           : logit

                                (Std. Err. adjusted for 76 clusters in proj_schid)
------------------------------------------------------------------------------
             |               Robust
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
absence17    |
    baseline |    3.5596   .5420783     6.57   0.000     2.497146    4.622054
      gender |  -.0092239    .051126    -0.18   0.857     -.109429    .0909812
         fsm |   .0414224   .0720544     0.57   0.565    -.0998016    .1826464
         sen |  -.1334815    .089411    -1.49   0.135    -.3087238    .0417609
         ks1 |  -.0244722   .0104919    -2.33   0.020     -.045036   -.0039084
      trial2 |  -.6563911   .0723573    -9.07   0.000    -.7982088   -.5145735
  condpr1_tot |    .033732   .0188903     1.79   0.074    -.0032922    .0707563
concentration1|  -.0475409   .0316891    -1.50   0.134    -.1096504    .0145686
   prosocial1 |  -.0512808   .0419448    -1.22   0.221    -.1334911    .0309296
  school_size |  -.0001126   .0001905    -0.59   0.554     -.000486    .0002607
   school_fsm |   .0085665   .0034252     2.50   0.012     .0018532    .0152798
  school_lev4 |   .0017915   .0035492     0.50   0.614    -.0051648    .0087479
       schcond|  -.0541339   .0665529    -0.81   0.416    -.1845752    .0763073
   school_eal |  -.0006847    .00157    -0.44   0.663    -.0037619    .0023925
        _cons |  -2.22828   .5622886    -3.96   0.000    -3.330345   -1.126214
ln(sessions17)|          1  (exposure)
-------------+----------------------------------------------------------------
c50          |
        _cons |         15  (constrained)
------------------------------------------------------------------------------
```

- The effect of **trial** for compliers *-0.656* (IRR=0.519)
- This means that compliers in the GBG have an incidence rate of 51.9% that of would-be compliers in the control group
- The GBG seems to be successful in **reducing absence** for those who played the game for sufficiently long

Final remarks

gsem is a flexible suite of Stata commands that allows us to estimate complex statistical models to assess trial efficacy under noncompliance (and much more).

We can make use of various **gsem** specifications to address different issues:

- Various distributional assumptions for outcome measures
- Standard error options for clustering
- Use of constraints to test different hypotheses (common coefficients and/or variances)
- Model comparison and selection using **information criteria** (AIC and BIC available)
- Fast, easy to use and widely available

Thank you!

Acknowledgements

This work was supported by the funding of:

- Education Endowment Foundation (no grant number)
- National Institute for Health Research (grant number 14/52/38)

Data access:

The Department for Education (DfE) granted access to the National Pupil Database to analyse the educational outcomes of participant children in the Secure Research Service, hosted by the UK Office for National Statistics.

Disclaimer:

This work was produced using statistical data from the **UK Office for National Statistics (ONS)**. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

References and further reading

Humphrey, N., et al. (2018). Good Behaviour Game: Evaluation report and executive summary. EEF. available here

Humphrey, N., Hennessey, A., Troncoso, P., Panayiotou, M., et al. (in press). Examining the impact of the Good Behaviour Game on health- and educational-related outcomes for children: a cluster RCT and cost-consequence analysis. Public Health Research.

Peugh, J., Strotman, D., McGrady, M., Rausch, J., and Kashikar-Zuck, S. (2017). Beyond intent to treat (ITT): A complier average causal effect (CACE) estimation primer, Journal of School Psychology, Volume 60, 2017, Pages 7-24. available here

Rabe-Hesketh, S., Skrondal, A., and Pickles, A. (2004). Generalized multilevel structural equation modeling. Psychometrika 69: 167–190. available here

Skrondal A., and Rabe-Hesketh, S. (2004). Generalized latent variable modeling: multilevel, longitudinal and structural equation models. London: Chapman & Hall.

Troncoso, P., and Humphrey, N. (forthcoming 2021). Playing the Long Game: A Multivariate Multilevel Non-Linear Growth Curve Model of Long-Term Effects in a Randomized Trial of the Good Behavior Game. Journal of School Psychology.