Causal inference in randomised trials

Dr Richard Emsley
Centre for Biostatistics, Institute of Population Health,
The University of Manchester, Manchester Academic Health Science Centre

MRC North West Hub for Trials Methodology Research

Visiting Lecturer, Institute of Psychiatry, Psychology and Neuroscience, KCL

http://www.population-health.manchester.ac.uk/staff/RichardEmsley/
richard.emsley@manchester.ac.uk

Australian Pharmaceutical Biostatisticians Group, Sydney
Monday 7th December 2015

Research Programme: Efficacy and Mechanisms Evaluation

Joint work with **Graham Dunn**, **Ian White, Andrew Pickles** and **Sabine Landau.**

Funded by Medical Research Council Methodology Research Programmes:

- **Design and methods of explanatory (causal) analysis for randomised trials of complex interventions in mental health (2006-2009)**
  - Graham Dunn (PI), Richard Emsley, et al
- **Estimation of causal effects of complex interventions in longitudinal studies with intermediate variables (2009-2012)**
  - Richard Emsley (PI), Graham Dunn.
- **MRC Early Career Centenary Award (2012-13)**
- **Designs and analysis for the evaluation and validation of social and psychological markers in randomised trials of complex interventions in mental health (2010-12)**
  - Graham Dunn (PI), Richard Emsley, et al.
- **Developing methods for understanding mechanism in complex interventions (2013-16)**
  - Sabine Landau (PI), Richard Emsley, et al.
- **MRC NorthWest Hub for Trials Methodology Research (2013-2018)**
  - Paula Williamson (PI), Richard Emsley, et al.

The four key questions about treatments

1. Does it work?
   - Efficacy analysis
2. How does it work?
   - Mediation analysis
3. Who does it work for?
   - Stratified/personalised medicine
4. What factors make it work better?
   - Process evaluation

CAUSAL INFERENCE AND RANDOMISED TRIALS

Need for a framework

- We want to assess causal effects of treatments on clinical outcomes from study data.

- To be able to do this we need a framework that:

  a. clearly defines the effects (population parameters) that we are interested in and

  b. allows us to work out when (under what assumptions) a statistical method provides an unbiased estimate of this parameter.

- This is the key insight of causal inference.

Potential outcomes approach

- The potential outcomes approach
  - It is a comparison between what is and what might have been.
  - It is counterfactual.
  - We wish to estimate the difference between a patient's observed outcome and the outcome that would have been observed if, contrary to fact, the patient's treatment or care had been different (Neyman, 1923; Rubin, 1974).

- We need the possibility of comparison to define an individual treatment effect.
  - (Note that a potential outcomes approach cannot be used to define causal effects of exposures that cannot be manipulated, such as gender.)

Defining treatment effects on outcome

- Consider a randomised controlled trial with two arms: treatment (R_i=1) versus control (R_i=0) and a continuous outcome Y

- Prior to randomisation to one of two competing treatment arms we can envisage two potential outcomes for each participant in the trial:
  - the outcome after receiving treatment, $Y_i(1)$
  - the outcome after receiving the control, $Y_i(0)$

- For a given individual, the effect of treatment is the difference:

$$\text{ITE}(Y)=Y_i(1)-Y_i(0)$$

- The observed outcome is:

$$Y_i = R_iY_i(1)+(1-R_i)Y_i(0)$$
$$= Y_i(0) + R_i\,(Y_i(1)-Y_i(0))$$

Average treatment effects

- The average treatment effect is given by:

$$\begin{aligned} \textbf{ATE} &= \textbf{Ave(ITE)} \\ &= \textbf{Ave}[\boldsymbol{Y}\textbf{(1)} - \boldsymbol{Y}\textbf{(0)}] \\ &= \textbf{Ave}[\boldsymbol{Y}\textbf{(1)}] - \textbf{Ave}[\boldsymbol{Y}\textbf{(0)}] \end{aligned}$$

- **If** allocation to treatment is **random**, and there is perfect compliance with the allocation, then

$$= \textbf{Ave}[\boldsymbol{Y}\textbf{(1)}|\boldsymbol{R}\textbf{=1}] - \textbf{Ave}[\boldsymbol{Y}\textbf{(0)}|\boldsymbol{R}\textbf{=0}]$$

$$= \textbf{Ave}[\boldsymbol{Y}|\boldsymbol{R}\textbf{=1}] - \textbf{Ave}[\boldsymbol{Y}|\boldsymbol{R}\textbf{=0}]$$

- This can be estimated by the difference between the mean outcome for those receiving treatment and the mean outcome for those in the control condition.

When can ATE be estimated?

- We wish to evaluate the effects of receiving a treatment compared to a suitably defined control condition.

- When can we do this without running into problems?
  - Randomised controlled trial (RCT): Participants are randomised to two arms (experimental treatment and control)
    - YES - provided participants adhere to their allocated treatment.
  - Observational study: Compares subjects receiving the experimental treatment with subjects under the control condition
    - SOMETIMES - only if variables that drive treatment group selection have been measured and accounted for appropriately.

A ‘perfect’ randomised controlled trial

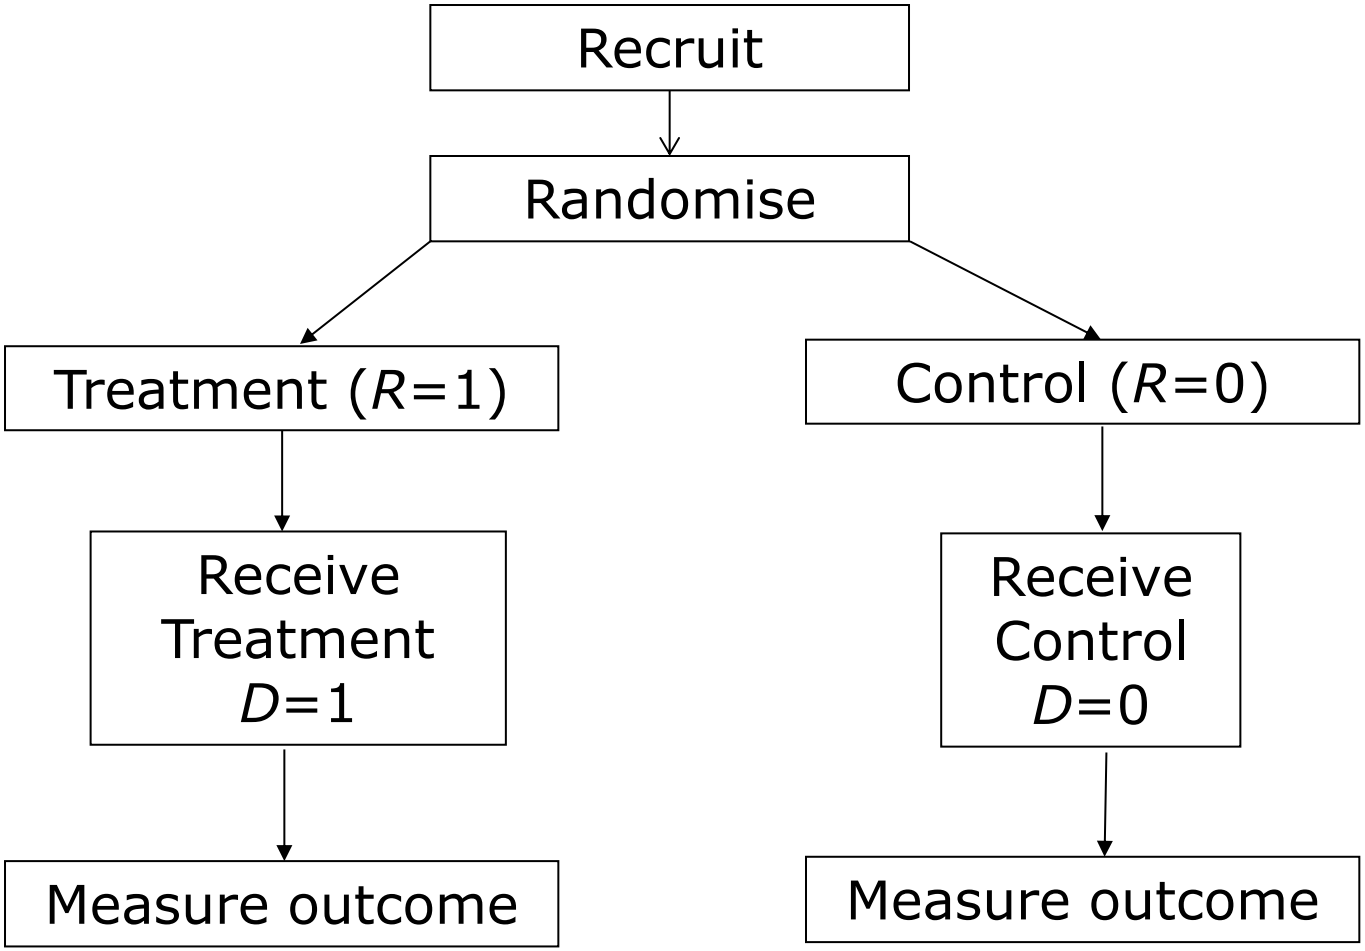

What are we estimating in trials?

- Interested in various measures of effect
  - Effectiveness - the benefit of a treatment policy
  - Efficacy - the benefit of actually receiving treatment
- ITT measures effectiveness as **implemented in a given trial**
- **What is the effectiveness of offering the intervention?**
- It tells us whether randomising the treatment works
  - **On average, not for an individual patient!**
  - **Regardless of whether you receive the treatment or not!**

What does intention-to-treat estimate?

- In trials with full compliance:
  - effectiveness = efficacy
  - ITT = ATE
- But in the presence of non-compliance in trials:
  - effectiveness ≠ efficacy
  - ITT analysis cannot assess efficacy (ATE)
- As a patient, more interested in efficacy?
- ITT is (probably) the parameter of interest for:
  - policy interventions
  - Health care providers
  - Health economic evaluation

Adjusting for baseline with continuous outcomes

- When baseline values of the outcome variable have been measured alternative estimators are available:
  - Difference in mean post treatment outcomes between trial arms (post estimator)
  - Difference in change scores (over time) between trial arms (change score estimator)
  - Conditioning on baseline value (ANCOVA estimator).

- All are known to be unbiased for ATE in trials; ANCOVA estimator can be shown to be most precise.

- For any conditioning variable X, assuming no interaction,

$$E[Y(1)-Y(0) \mid X] = ATE=ATE|X$$

Non-collapsibility of the OR

- For binary clinical outcomes non-linear models are more appropriate to describe the relationship between treatment and the expected outcome.

- Unfortunately even in the absence of treatment effect modification by baseline variable X, conditional and marginal ORs are not the same: $\exp(\Phi) \neq \exp(\Phi|X)$

- This is known as non-collapsibility of the OR (for more details see Hernan *et al.*, 2011)
  - Risk Ratio and Risk Difference are collapsable

- As a consequence an OR estimator will estimate a different treatment effect estimand depending on the covariates conditioned on in the analysis model.

Missing outcomes

- Most analysis valid under MAR (White & Thompson, 2005)
- MAR will depend on the covariates in the analysis model
  - Implications for binary outcomes?
  - What happens when there is non-compliance with randomisation in an ITT analysis?

ADJUSTING FOR DEPARTURES FROM RANDOM ALLOCATION: BASIC CONCEPTS

An ‘imperfect’ RCT

There could not be worse experimental animals on earth than human beings; they complain, they go on vacations, they take things they are not supposed to take, they lead incredibly complicated lives, and, sometimes, they do not take their Medicine.

Efron B. Foreword. Statistics in Medicine
1998; 17: 249–50.

Treatment Allocation and Treatment Receipt

- Subjects are randomised to receive treatment (R_i=1) or not (R_i=0).
- As for outcomes, prior to treatment allocation, there are two **potential treatments** for each subject:
  - $D_i(R_i = 1)$ or $D_i(1)$
  - $D_i(R_i = 0)$ or $D_i(0)$
- If allocated to receive treatment and complies with the allocation then $D_i(1) = 1$.
  - Otherwise $D_i(1) = 0$.
- If allocated to the control condition but actually receives treatment then $D_i(0) = 1$,
  - otherwise $D_i(0) = 0$.

A more realistic RCT

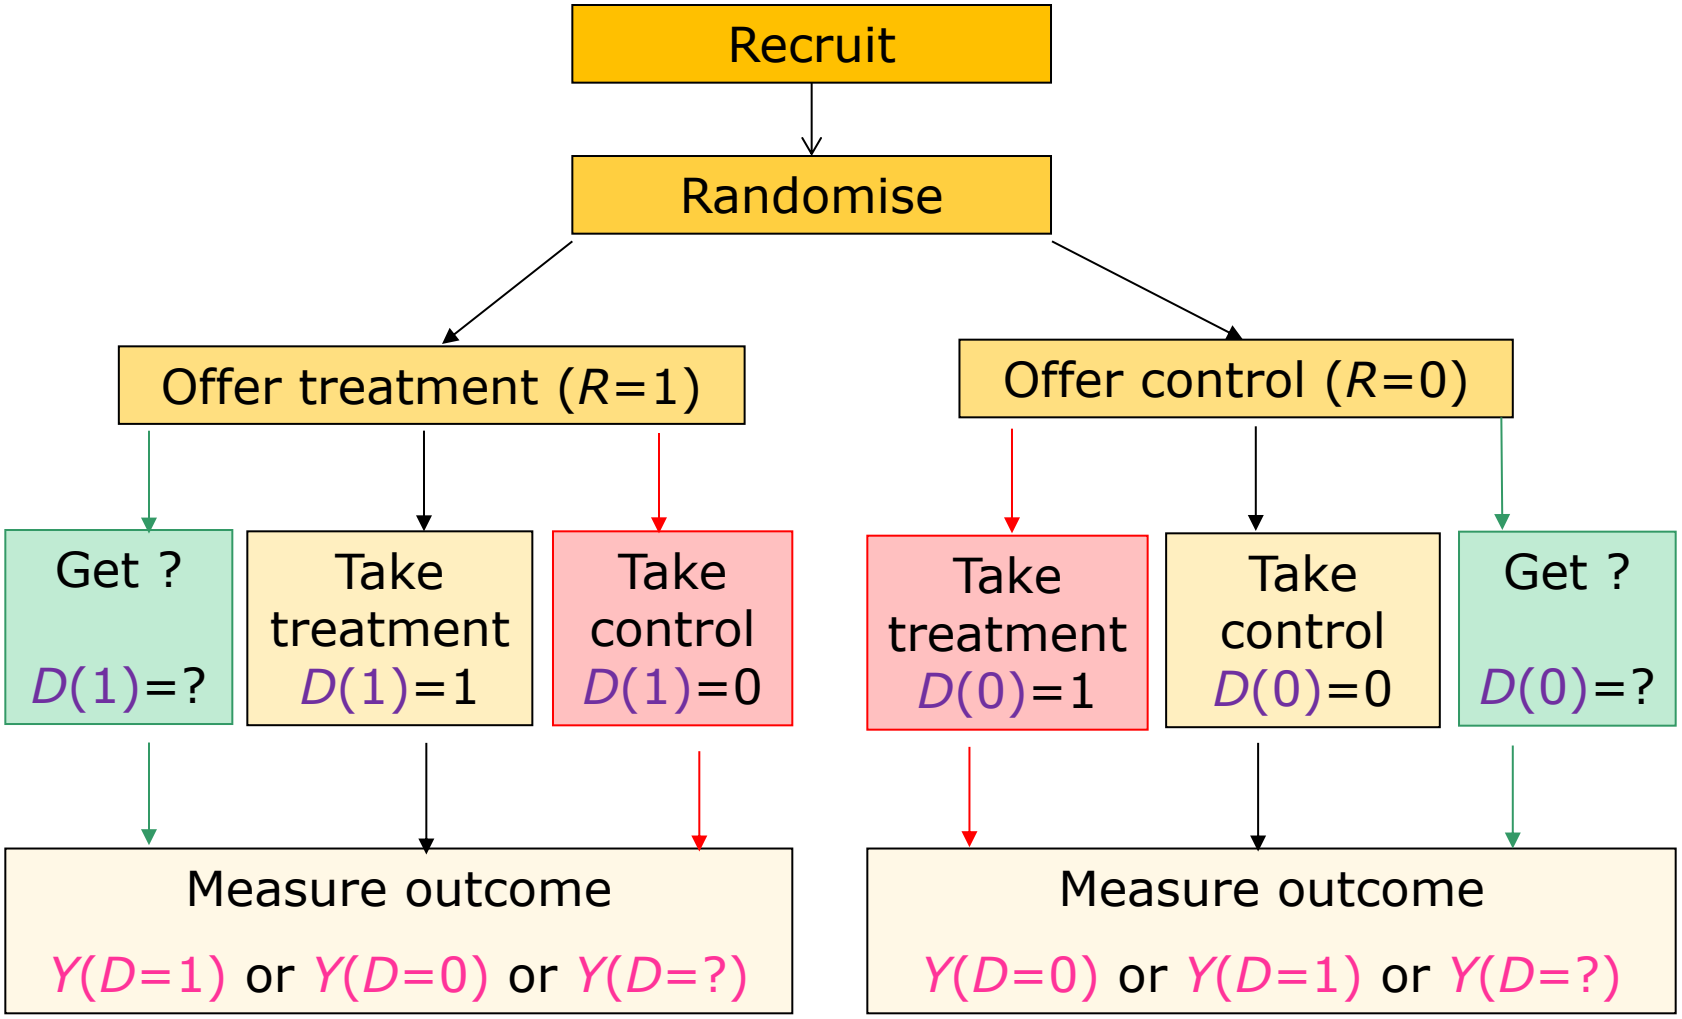

Commonly-used approaches

1. ***ITT*** – estimates the effect of offering treatment.
   - Valid if the effect of offering is the same as the effect of receiving treatment.

2. ***Per Protocol*** – compares those who complied with the treatment offer will ***all*** of the controls.

3. ***As Treated*** – compares those who received treatment with everyone who did not (i.e. abandons randomisation).

Special case: no treatment access

- Consider the case where those randomised to control have no access to treatments other than those forming part of the control condition
- Those randomised to receive the experimental treatment receive the control condition if they do not comply.
- Then the RCT simplifies to

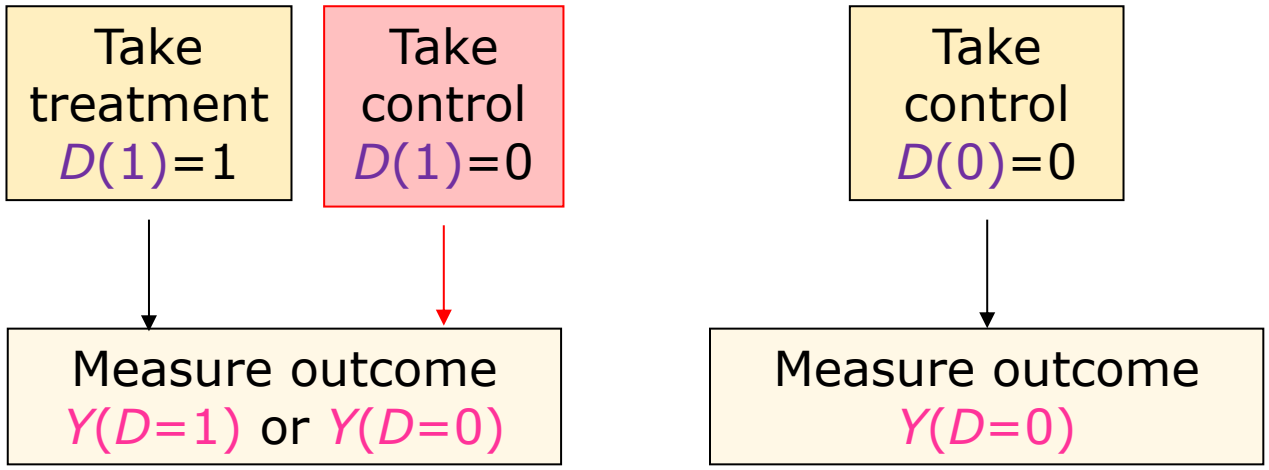

Per-protocol approach

- Per-protocol analysis = comparison of outcome in those who take the treatment versus outcome in all those randomised to control.
- Non-compliers in the active treatment group are excluded from the analysis.

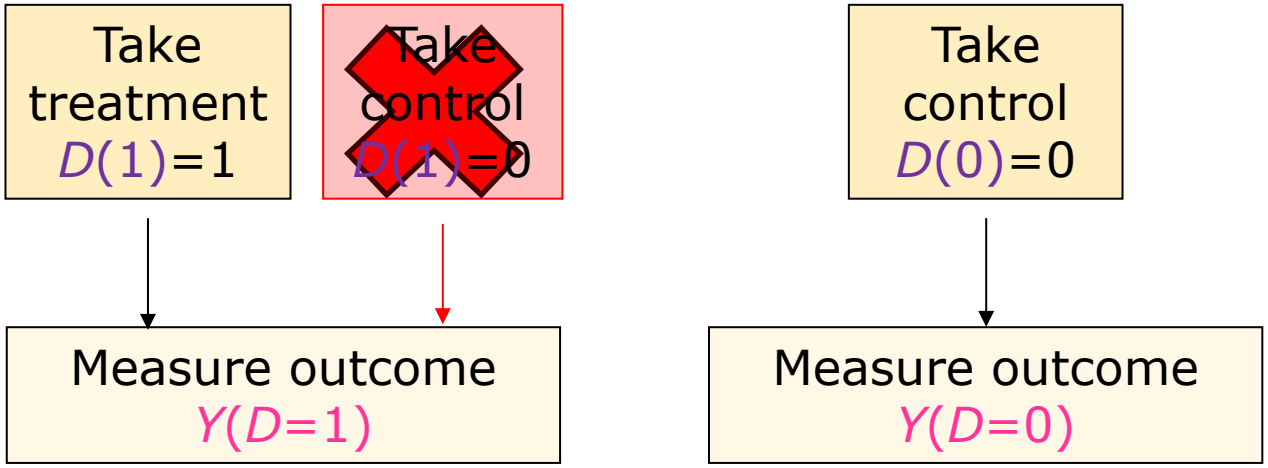

As-treated approach

- Per-treated analysis = comparison of outcome in those who take the treatment versus outcome in those who take control, regardless of randomisation.

- Non-compliers in the active treatment "moved" into the control group.

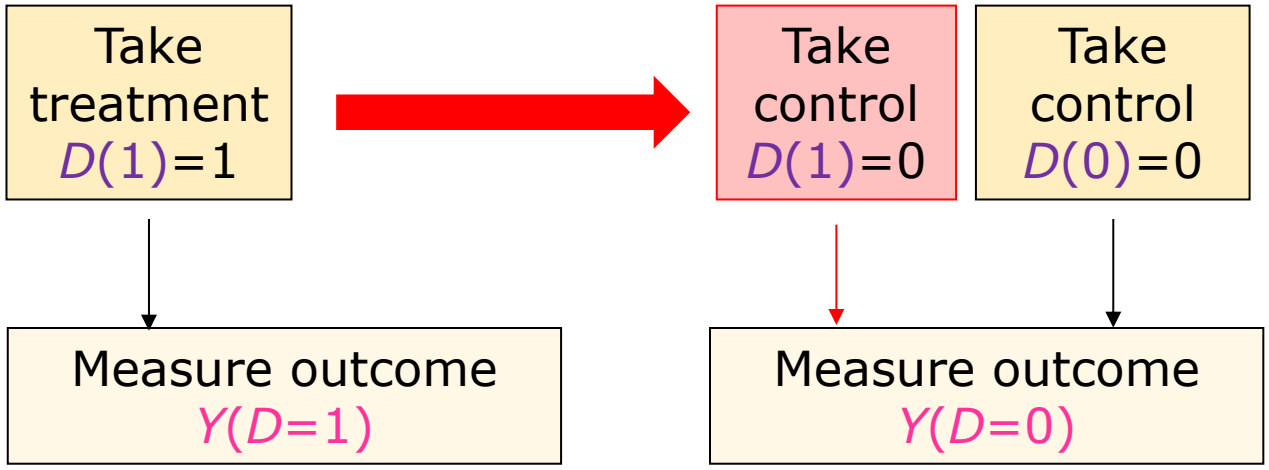

The problem of confounding

- The difference Ave[$Y|D=1$] − Ave[$Y|D=0$] is usually ***not*** a valid (unbiased) measure of the average causal effect.
- We infer that there are variables (confounders) which account for these biases. They can be either measured (X) or unmeasured (U).
- No confounding for the average causal effect if both

$$\text{Ave}[Y(0)|D=1] = \text{Ave}[Y(0)|D=0]$$

$$\text{Ave}[Y(1)|D=1] = \text{Ave}[Y(1)|D=0]$$

- In words, the mean of potential outcomes for the control condition is not dependent on whether the participant actually receives treatment.
- Similarly the potential outcomes for the treatment condition are not influenced by treatment actually received.

The Complier Average Causal Effect (CACE)

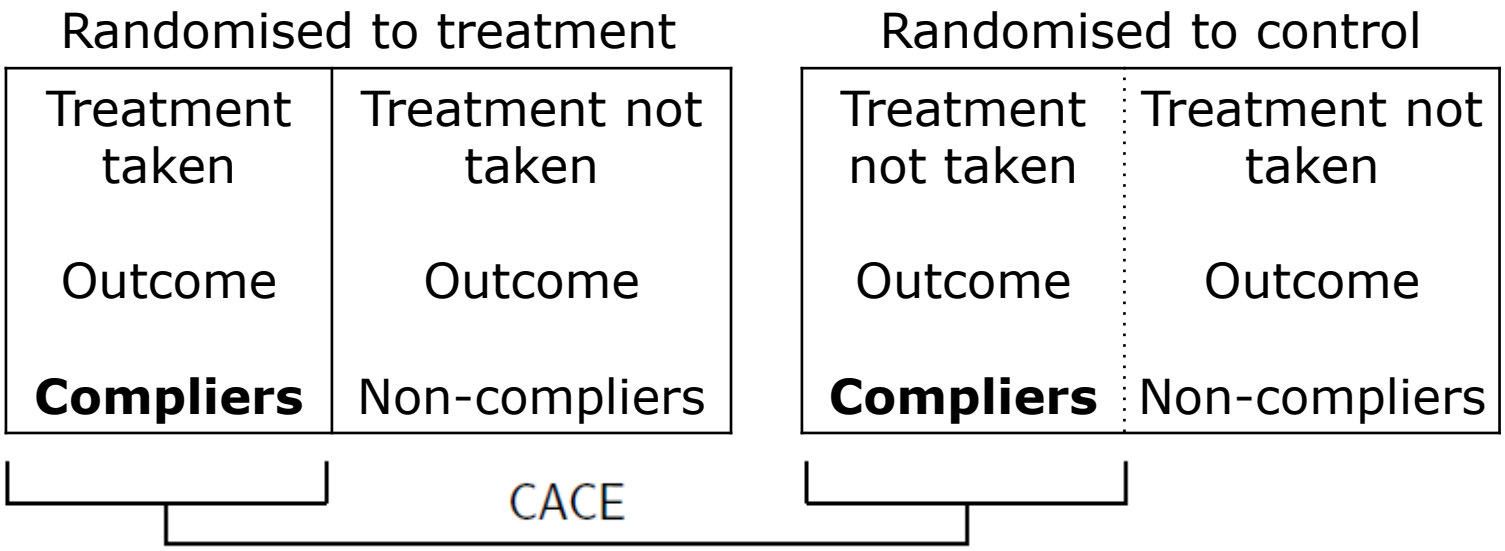

- The ***Complier-Average Causal Effect (CACE) estimate*** is the comparison of the average outcome of the compliers in the treatment arm with the average outcome of the comparable group of would-be compliers in the control arm.
- This is a randomisation-respecting estimate.
- It is the ITT effect in the sub-group of participants who would always comply with their treatment allocation. It is not subject to confounding.

Average Treatment Effects

Average Treatment Effect (*ATE*):

$\boldsymbol{\Delta} = \text{Ave}[Y(1)-Y(0)]$

Average Treatment effect on the Treated (*ATT*):

$\boldsymbol{\Delta}_{\textbf{Treated}} = \text{Ave}[(Y(1)-Y(0)) \mid D=1]$

Complier Average Causal Effect (*CACE*)*:

$\boldsymbol{\Delta}_{\textbf{Complier}} = \text{Ave}[(Y(1)-Y(0)) \mid (D(1)-D(0)=1)]$

*or Local Average Treatment Effect (*LATE*)

Efficacy (CACE) estimation: assumptions

1. There are two latent classes of participants: Compliers and Non-compliers. Compliers get treatment if and only if allocated to the treatment condition. Non-compliers never get the treatment, regardless of allocation.
   - We can identify these two groups in the treatment arm, but they remain hidden (unobserved) in the control arm.

2. As a consequence of randomisation, on average, the proportion of Compliers is the same in the two arms of the trial.

3. In the absence of treatment (i.e. for the Non-compliers) there is no effect of randomisation (i.e. treatment arm) on outcome. This assumption is often called an exclusion restriction.

A hypothetical RCT with non-compliance

- Size of trial = 1000 participants (500 in each group)
- 500 are allocated to receive chemotherapy alone.
- 500 are allocated to receive chemotherapy plus radiotherapy (the latter, say, starting a couple of months into a 6 month treatment period).
- Of the latter, for whatever reason (but likely to be linked with prognosis or initial progress under chemotherapy) only 250 actually get their allocated radiotherapy.
- Outcome? Tumour volume at 6 months after randomisation.

The results of our hypothetical RCT

# participants Mean tumour volume (ml)	Compliers	Non-compliers	All
Radiotherapy + Chemo	250 (50%) 40	250 60	500 50
Chemo alone	? ?	? ?	500 100

CACE estimation

# participants Mean tumour volume	Compliers	Non-compliers	All
Radiotherapy + Chemo	250 (50%) 40	250 60	500 50
Chemo alone	*250* *140*	*250* *60*	500 100

Complier-Average Causal Effect(CACE)

randomisation balance

exclusion restriction

CACE = 40 – 140 = -100
(cf ITT = 50 – 100 = -50)

Comparison of Estimates

ITT:	50 - 100	=	-50
CACE:	40 - 140	=	-100
Per Protocol:	40 - 100	=	-60
As Treated:	40 – 87	=	-47

Per Protocol?

- Based on the difference

  $$\text{Ave}[Y(1)|R=1 \ \& \ \text{Complier}] - \ \text{Ave}[Y(0)|R=0]$$

- ***Only*** valid if

  $$\text{Ave}[Y(0)|\text{Complier}] \ = \ \text{Ave}[Y(0)|\text{Non-complier}]$$

- i.e. implicitly based on the assumption that the treatment-free prognosis is independent of compliance class.

- Otherwise, Per Protocol estimates are subject to confounding (selection effects).

As Treated?

- Based on the difference

  Ave[$Y(1)$|Treated] - Ave[$Y(0)$|Not treated]

- ***Only*** valid if the exclusion restriction holds **and**

  Ave[$Y(0)$|Complier] = Ave[$Y(0)$|Non-complier]

- This is very unlikely!

CACE and assumptions

- The estimators are valid if and only if the assumptions are true.
- Which assumptions seem to be more realistic?
- The vital assumption for the CACE estimator is the exclusion restriction (no effect of randomisation on outcome except through treatment received).
- The exclusion restriction is likely to hold in a double-blind placebo-controlled drug trial, but what about the case of an unblinded trial of psychotherapy?
  - Resentful demoralisation.

Problems with the CACE

- We need to be able to define (binary) compliance.
- We don't really know who is a "Complier", particularly in the control group.
- In practice, we may want to know what will be observed…
  - …if compliance is worse than in the trial (e.g. if rolled out in clinical practice)
  - …if compliance is better than in the trial (e.g. because intervention is well publicised / marketed)
- This means we may want to know the average causal effect in a different subgroup. We might assume this is the CACE – but it is an assumption.

Formal CACE estimation

- So how can we construct estimates of CACE and associated standard errors?

- Possible estimation methods:
  - Principal stratification
  - "By hand" approach
  - Negative weights
  - Instrumental variables approaches
  - Inverse probability weighting

CACE estimation "by hand"

- We have already seen that when we have only **Compliers** and **Non-compliers** then CACE can be estimated by a set of simple operations based on information from the three observed strata.

- To get an even quicker answer under this scenario consider the effect of the random treatment offer:

  $ITT = p_C * ITT_{Compliers} + (1-p_C) * ITT_{Non\text{-}compliers}$

  $= p_C * ITT_{Compliers}$ (from the exclusion restriction)

- CACE estimate = estimate of $ITT_{Compliers}$

  = estimate of ITT / estimate of p_C

  = -50/0.5 = -100

The instrumental variable model

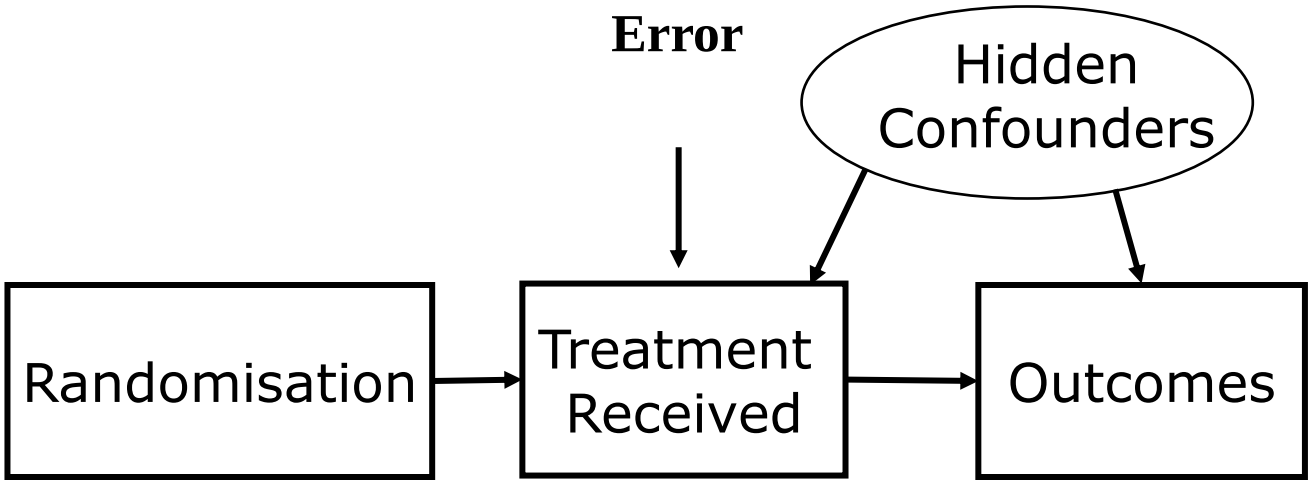

- Randomisation influences both treatment receipt and outcome. However, the influence on outcome is only through treatment received. Confounders are independent of randomisation.

- By assuming the absence of a direct path from randomisation to outcome, we assume the entire effect of randomisation acts through receipt of treatment.
  - **randomisation is an instrumental variable.**

CACE and principal stratification: allowing for contamination/crossover

- Consider our counterfactual treatment received variables:
  - $D_i(1)$= treatment if randomised to intervention
  - $D_i(0)$= treatment if randomised to control
  - both are 0/1 (received standard / intervention)

- Implies 4 types of person (compliance-types here):
  - $D_i(1)=1$, $D_i(0)=1$: always-takers
  - $D_i(1)=1$, $D_i(0)=0$: compliers
  - $D_i(1)=0$, $D_i(0)=0$: never-takers
  - $D_i(1)=0$, $D_i(0)=1$: defiers – usually assumed to be absent

- Principal stratification was an idea of Frangakis and Rubin (1999), generalising the simple compliance-types above, where principal strata are the levels of the pair ($D_i(1)$, $D_i(0)$)

CACE and principal stratification

Class *C* (stratum)	$D_i(1)$	$D_i(0)$	$(D_i(1),D_i(0))$	Proportion	Treatment Effect
Never-takers	0	0	(0,0)	p_1	ITT_{never}
Always-takers	1	1	(1,1)	p_2	ITT_{always}
Compliers	1	0	(1,0)	p_3	$ITT_{compliers}$
Defiers	0	1	(0,1)	p_4	$ITT_{defiers}$

- There are now four distinct possibilities (classes) for the joint combination of $D_i(1)$ and $D_i(0)$.
- It is sometimes assumed that only one of classes 3 and 4 is present (known as the monotonicity assumption).

Using principal stratification approach to CACE

- We should model outcomes conditional on principal strata
  - allow differences between randomised groups within principal strata
  - these parameters have a causal meaning

- However, this is complicated since individual class membership is not known: for example, an individual with R_i=1 and D_i=1 is only known to belong to always-takers or compliers.

- We learn nothing about the effect of treatment from the always and never takers, since their treatment doesn't change.

- With these extra classes, this can be modelled using a finite mixture model.

INTRODUCTION TO MEDIATION ANALYSIS

Mediation and mediators

- Baron and Kenny (1986) defined mediation as the "*generative mechanisms through which the focal independent variable is able to influence the dependent variable of interest*"
- A mediator (*M*) is a variable that occurs in the causal pathway from an exposure (*R*) to an outcome variable (*Y*). It causes variation in the outcome and itself is caused to vary by the exposure variable.
  - This causal chain implies a temporal relation
    - *R* occurs before *M* and
    - *M* occurs before *Y*
- Mediating variables are often called intervening or intermediate variables.
  - (They have also been called process variables; but we reserve this term for variables that measures aspects of the therapeutic process.)

Complete mediation

- To reflect the mediated effect we need a path from *R* to *M* and a further path from *M* to *Y*.
- The following diagram illustrates complete mediation by *M*.

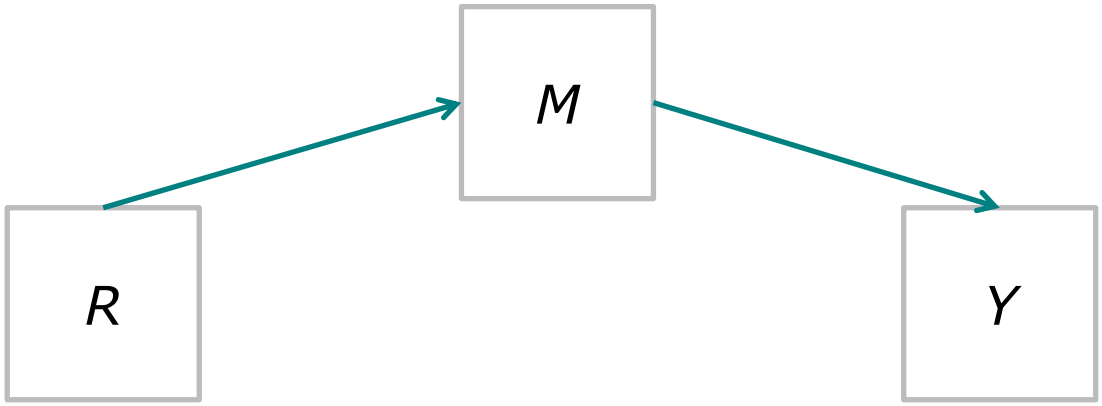

- Note that the diagram implies that *M* is the only mechanism by which *R* can change *Y.*

"The mediation triangle"

- We might not want to rule out effects of *R* on *Y* other than those operating by changing *M*.

- The following triangle illustrates partial mediation by *M*.

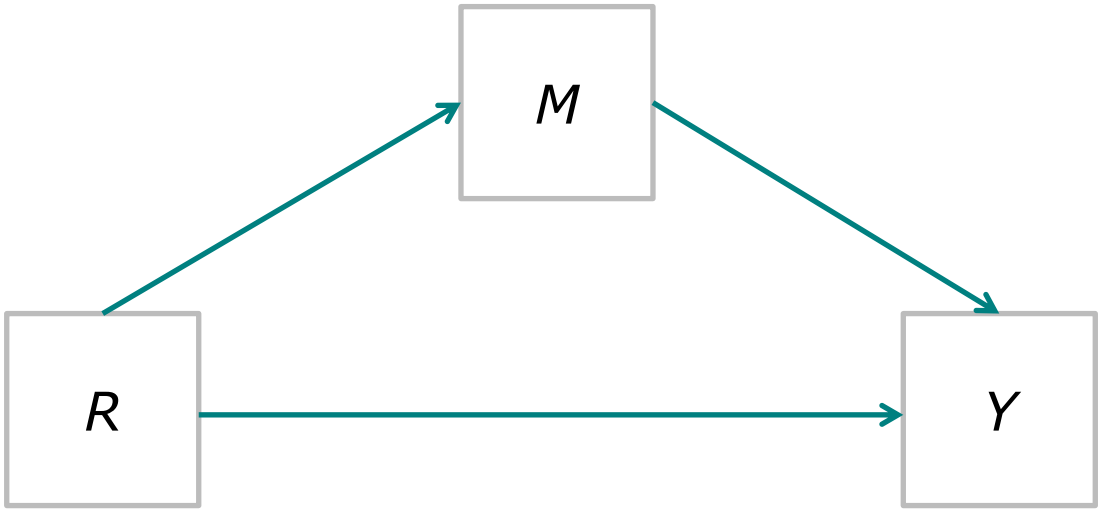

Mediation example

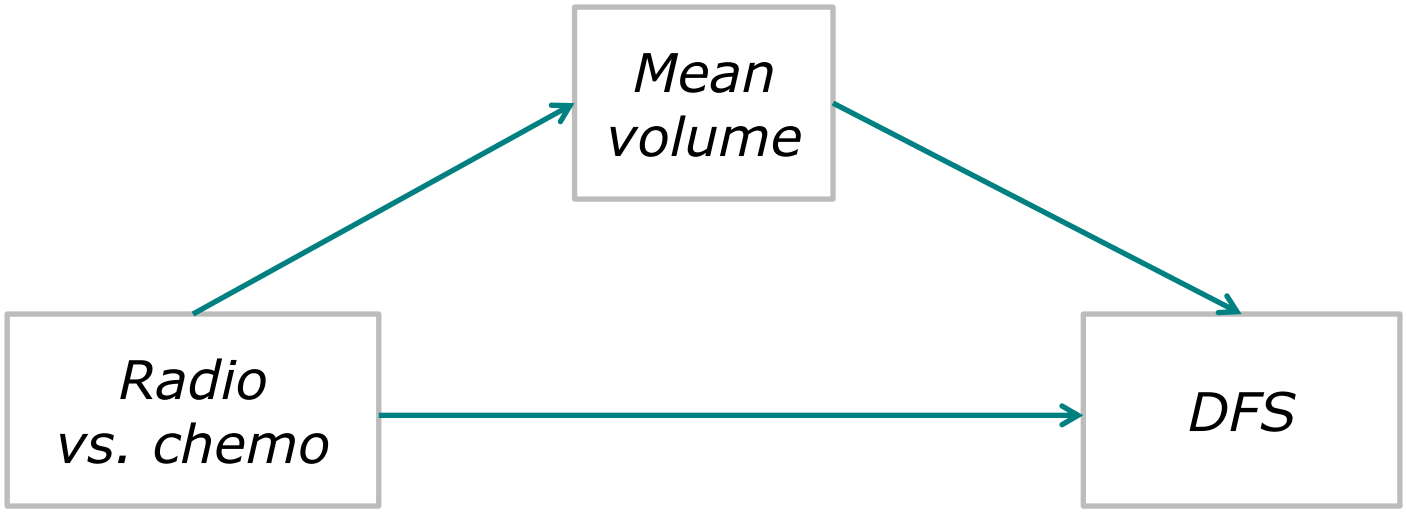

Direct and indirect effects

- Mediation investigations in trials aim to partition total (causal) treatment effects into
  1. effects that operate via changing the putative mediator – so called indirect treatment effects
  2. and non-mediated effects – so-called direct treatment effects.
- **Total effect = direct effect + indirect effect**
- Note that direct effects include effects via any mediating variable not included in the model.
  - So the meaning of a direct effect is always relative to the variable whose mediating effect is being modelled.

The basic underlying problem: estimating valid causal effects

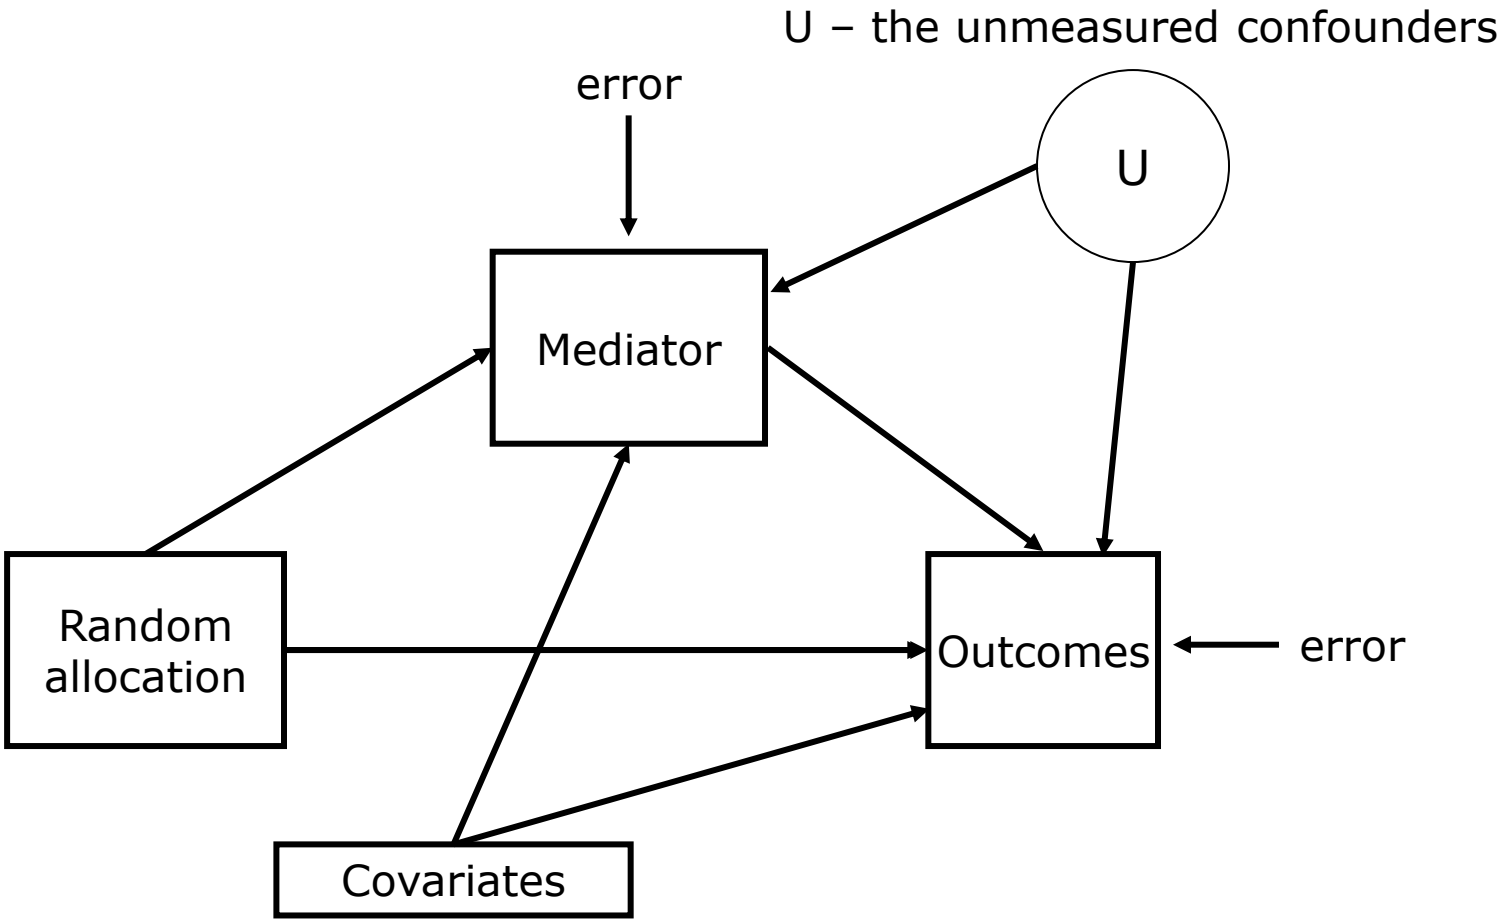

Traditional regression approach

- The traditional regression approach popular in the social and behavioural sciences was first mentioned in Hyman (1955), and further proposed by Judd and Kenny (1981) and Baron and Kenny (1986).
- It is based on two regression models:
  - Model for mediator (*M*): $E[M|R,X] = \mu + \alpha R + \delta X$
  - Model for outcome (*Y*): $E[Y|R,M,X] = \tau + \gamma R + \beta M + \varphi X$
  - (**X** are baseline covariates that act as observed confounders.)
  - γ is the direct effect of *R* on *Y* (not through *M*)
  - $\alpha \times \beta$ is the indirect effect (through *M*)

Problems with traditional approaches

1. Definitions are specific to linear models
2. Assumes no interactions between *R*-*M*
3. Unmeasured confounding of *M*-*Y* relationship leads to bias in partitioning of direct and indirect effects
4. No measurement error in *M*
5. Intermediation confounding
   - Common causes of *M* and *Y* that are affected by *X*

Intermediate confounding

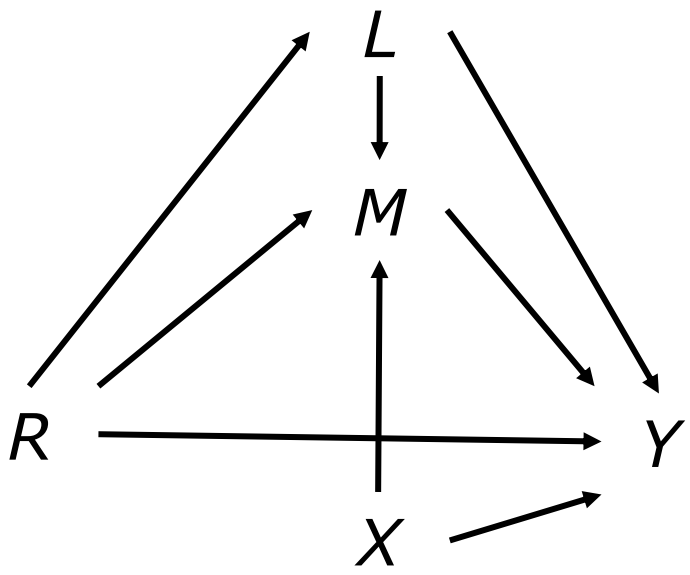

- *L* is a confounder for *M*-*Y*
- *L* is on the causal pathway from *R*-*Y*
- We both **should** and **should not** adjust for *L*

Causal mediation analysis

- Traditional approach has four main problems:
  1. Unmeasured confounding between mediator and outcome
  2. No interactions between exposure and mediator on outcome
  3. Doesn't easily extend to non-linear models
  4. Assumes correctly specified models

- Causal mediation analysis has arisen from the causal inference literature, and addressed these problems.

- Formally defines the causal mediation parameters.

- Extensively covered in VanderWeele (2015).

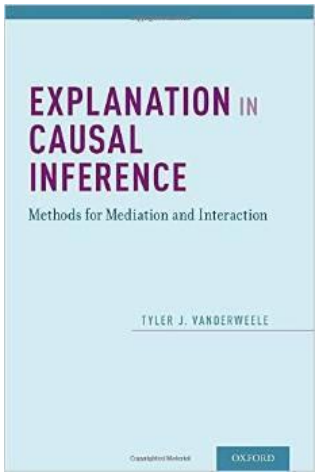

Causal mediation parameters

- Controlled Direct Effect:
  - the direct effect of treatment on outcome at mediator set to a particular value
  - **What is the effect of radiotherapy on survival if everyone in the population had the same tumour size?**

- Natural Direct Effect:
  - What is the effect of treatment when the mediator takes its "natural" level under the control condition?
  - **What is the effect of radiotherapy on survival if everyone has the tumour size they would have had on chemotherapy alone?**

- Natural Indirect Effect:
  - What is the effect of change in the mediator on clinical outcome if everyone is randomised to treatment?
  - **What is the effect of radiotherapy induced change in tumour size on survival if everyone received radiotherapy?**

Estimating causal parameters

- Wide range of options for most combinations of *M* and *Y*
- Work on identification and estimation of direct and indirect causal effects using parametric regression models
  - VanderWeele and Vansteelandt (2009, 2010):
  - Outcomes can be continuous, binary, count or survival.
  - Mediators can binary or continuous.
- G-computation is flexible and efficient but requires parametric modelling assumptions:
  - correct specification of all relevant conditional expectations and distributions
  - gformula command in Stata (Daniel et al., 2011)
- Semi-parametric methods make fewer parametric assumptions:
  - Inverse probability of treatment weighting (IPTW)
  - G-estimation

Incorporating compliance or departures from randomised treatment

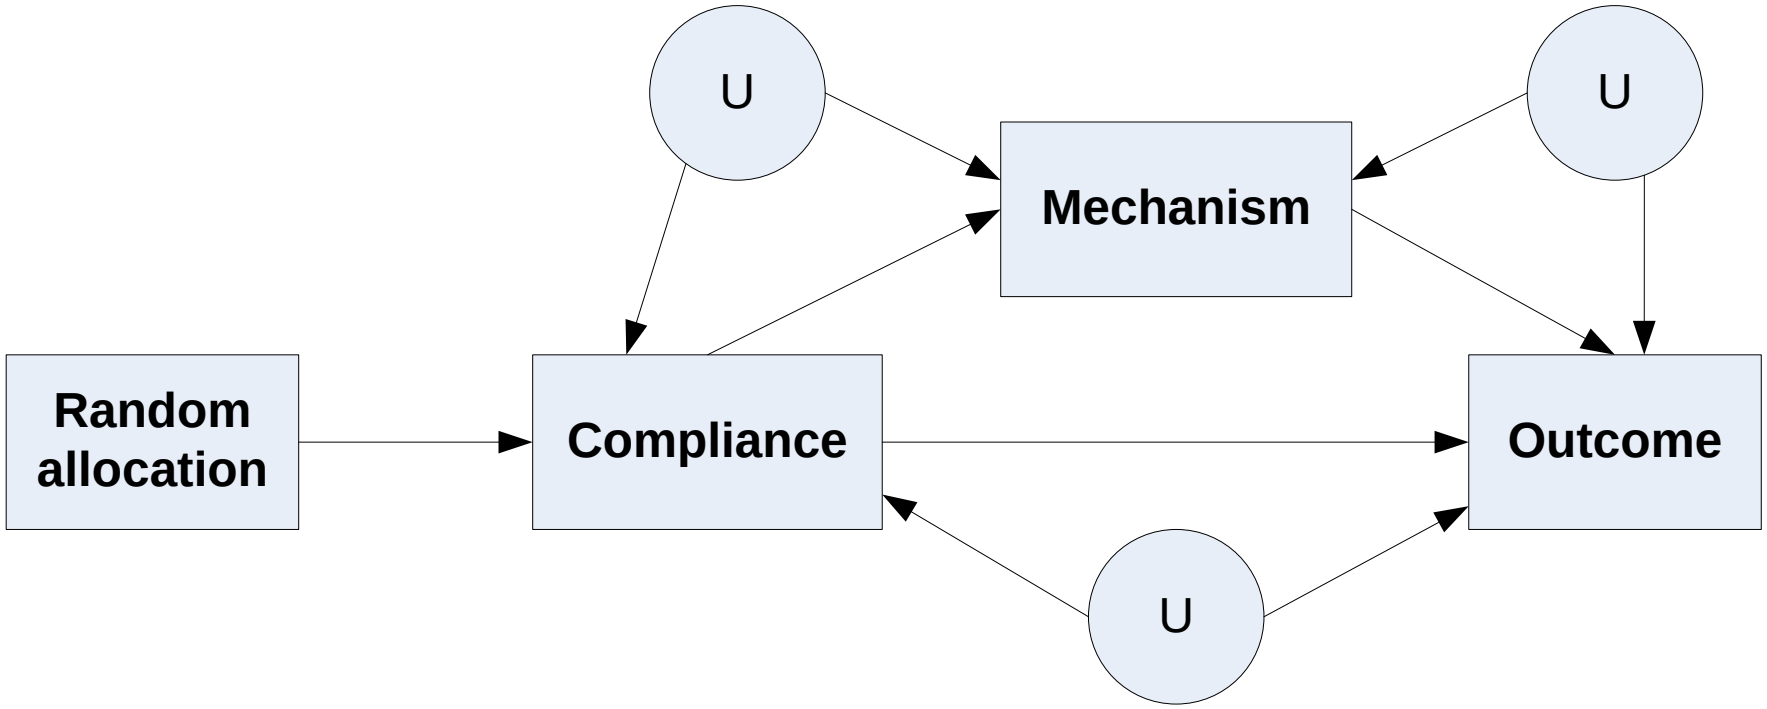

Summary: efficacy estimators

- In the presence of non-compliance with randomly allocated treatments ITT no longer assesses efficacy.
- We need to be clear about the subpopulation for which we estimate an average treatment effect.
- Commonly used efficacy estimators such as per-protocol and as-treated approaches attempt to estimate CACE.
- However, to do this they rely on the assumption of "no residual confounding".
- CACE can be estimated from trials without making this assumption.

Summary: mediation analysis

- In a RCT, if the treatment doesn't effect the mediator, then there can't be mediation through that variable.
  - Target effect
  - Only aspect which can be tested unbiasedly?

- Researchers in mediation potentially oblivious of implicit assumptions they are making but all solutions depend on assumptions.

- There are also lots of statistical challenges!

- Some of them have been solved but, there needs to be more research done on relevant trial designs (e.g. Imai et al, 2013).

Methodology report

- Dunn G, Emsley RA, Liu H, Landau S, Green J, White I and Pickles A. (2015). Evaluation and validation of social and psychological markers in randomised trials of complex interventions in mental health. Health Technology Assessment 19 (93).

- Non-technical introduction and summary of our work on analysing complex interventions:
  - Introduction to CI
  - Mediation analysis
  - Process evaluation
  - Longitudinal extensions
  - Stratified medicine
  - Guidance and tips for trialists

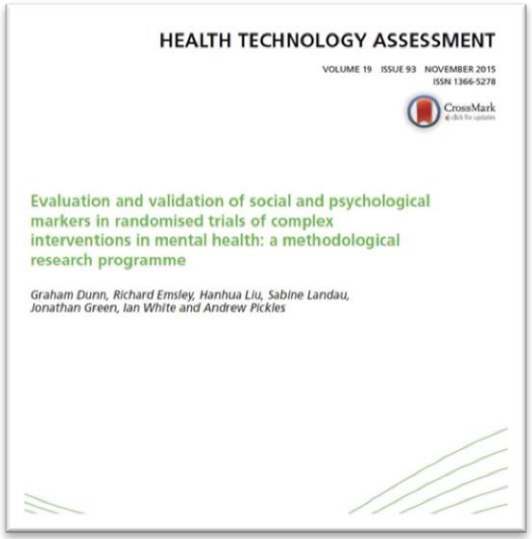

Some references (personal)

- Dunn G, **Emsley RA**, Liu H, Landau S, Green J, White I and Pickles A. (2015). Evaluation and validation of social and psychological markers in randomised trials of complex interventions in mental health. Health Technology Assessment 19(93).
- Dunn G, **Emsley RA**, Liu H & Landau S. (2013). Integrating biomarker information within trials to evaluate treatment mechanisms and efficacy for personalised medicine. *Clinical Trials*, 10(5):709-19.
- **Emsley RA** & Dunn G. Process evaluation using latent variables: applications and extensions of finite mixture models. (2013). In: *Advances in Latent Variables*. Eds. Brentari E., Carpita M., Vita e Pensiero, Milan, Italy. ISBN: 9788834325568.
- **Emsley RA** & Dunn G. (2012) Evaluation of potential mediators in randomized trials of complex interventions (psychotherapies). In: *Causal Inference: Statistical perspectives and applications.* Eds: Berzuini C, Dawid P & Bernardinelli, L. Wiley.
- **Emsley RA**, Green J, Dunn G. (2011). Designing trials of complex interventions for efficacy and mechanisms evaluation. *Trials*, 12(Suppl1), A143.
- **Emsley RA**, Dunn G & White IR. (2010). Modelling mediation and moderation of treatment effects in randomised controlled trials of complex interventions. *Statistical Methods in Medical Research*, 19(3), 237-270.
- VanderWeele TJ, **Emsley RA**. (2013). Discussion of "Experimental designs for identifying causal mechanisms". *JRSS-A*, 176(1), pp46.