MANCHESTER 1824

The University of Manchester

What is... Causal Inference?

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

www.medicine.manchester.ac.uk/staff/RichardEmsley
richard.emsley@manchester.ac.uk

Presentation for Methods@Manchester "What Is…" series
Thursday 14th February 2013

What is causal inference?

"First, for what reason we pronounce it *necessary*, that everything whose existence has a beginning, should also have a cause? Secondly, why we conclude, that such particular causes must *necessarily* have such particular effects; and what is the nature of that *inference* we draw from the one to the other, and of the belief we repose in it?"

David Hume, Treatise on Human Nature (1739-40), Book I (Of The Understanding), Part III (Of Knowledge and Probability)

"…statisticians must find a sound philosophical basis for causality and be able to express in some mathematical form the knowledge and assumptions needed for causal inference."

Stone (1993)

A brief history of causal inference (1)

- Neyman (1923) and Fisher (1925) discussed the **potential yield** to be gained from agricultural plots under different experimental exposures.

- First introduction of the concept of random allocation as an experimental design.

Ronald Fisher (1890-1962)

Jerzy Neyman (1894-1981)

A brief history of causal inference (2)

- This was formalised statistically for both randomised and non-randomised studies many years later.
  - Potential outcomes
  - Rubin Causal Model (Holland 1986)

- Rubin DB (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66(5), 688-701.
  - 1031 citations on Web of Science.

- Rosenbaum PR and Rubin DB (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1), pp41-55.
  - 3825 WoS citations.

Don Rubin

The general principle of causal inference

- We often compare an outcome after an action has occurred with what would have happened had the action not occurred, assuming all other things would have remained equal.
- This is often done mentally (and sometimes unconsciously).
- If the two outcomes differ, we say the action has had a *causal effect*.
- Can this be modelled statistically?

The general principle of causal inference

- Statistical models can only tell us about association between two variables (say *X* and *Y*).

- The aim of causal inference is to infer whether this association can be given a causal interpretation (e.g. *X* causes *Y*) by:
  - defining the causal estimands
  - being explicit about the assumptions being made
  - thinking about other possible explanations for observed effects, especially confounding.

- There are now many, many methods purporting to give causally valid solutions to this problem; this session only gives an overview of some of these.

How do we formally define a causal effect?

- Illustrated using the potential outcomes/counterfactual approach.
- It is a comparison between what is and what might have been.
- 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).
- Without the possibility of comparison the treatment effect is not well defined e.g. gender as a cause.

Causal inference is a comparison

Receive treatment

Measure outcome

Receive control

Measure outcome

Comparison of outcomes gives an
individual treatment effect

Causal inference is a comparison

Receive treatment

Measure outcome

Receive control

Measure outcome

Comparison of these outcomes **will not give** an individual treatment effect

Causal inference is a comparison

Receive treatment

Measure outcome

Receive control

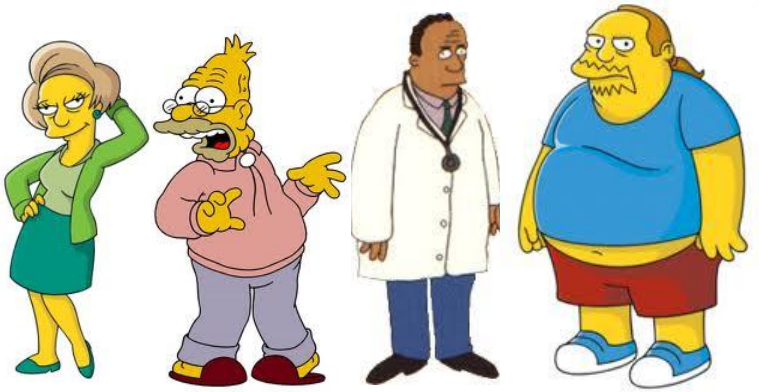

Measure outcome

Comparison of average outcomes gives an **average treatment effect**

Treatment effect heterogeneity

- The definition does not require that individual treatment effect is equal for everyone.

	Positive individual effect	Detrimental individual effect
Receive treatment		
Receive control		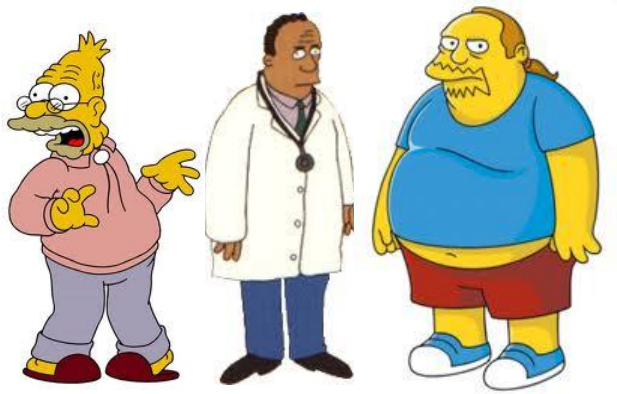

Individual treatment effects

- Consider a randomised controlled trial with two arms: treatment (Z=T) versus control (Z=C) and a continuous outcome Y
- Prior to randomisation we can envisage two potential outcomes for each participant in the trial:
  - the outcome after an active treatment, $Y(Z=\mathrm{T})=Y(\mathrm{T})$
  - the outcome after receiving the placebo, $Y(Z=\mathrm{C})=Y(\mathrm{C})$
- For a given individual, the effect of treatment is the difference:

$$\mathrm{ITE}(Y)= Y(\mathrm{T})-Y(\mathrm{C})$$

- As a result of the allocation, however, it is only ever possible to observe one of them (the other is a **counterfactual**).

Potential outcomes

Individual	Treatment Assignment		Individual effect $Y(T)-Y(C)$
	Treatment $Y(T)$	Control $Y(C)$	
1	6	3	3
2	4	6	-2
3	5	1	4
4	8	8	0
5	5	3	2
6	7	6	1
7	3	5	-2
8	9	4	5
9	4	2	2

Observed outcomes

Individual	Treatment Assignment		Individual effect $Y(T)-Y(C)$
	Treatment $Y(T)=Y$	Control $Y(C)=Y$	
1	6	-	?
2	-	6	?
3	5	-	?
4	8	-	?
5	-	3	?
6	7	6	?
7	-	5	?
8	9	-	?
9	4	-	?

The statistical solution - averages

The average treatment effect is given by:

$$\mathbf{ACE = Ave(ITE) = Ave[}\boldsymbol{Y}\mathbf{(T) - }\boldsymbol{Y}\mathbf{(C)]}$$
$$\mathbf{= Ave[}\boldsymbol{Y}\mathbf{(T)] - Ave[}\boldsymbol{Y}\mathbf{(C)]}$$

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

$$\mathbf{= Ave[}\boldsymbol{Y}\mathbf{(T)|}\boldsymbol{Z}\mathbf{=T] - Ave[}\boldsymbol{Y}\mathbf{(C)|}\boldsymbol{Z}\mathbf{=C]}$$

$$\mathbf{= Ave[}\boldsymbol{Y}\mathbf{|}\boldsymbol{Z}\mathbf{=T] - Ave[}\boldsymbol{Y}\mathbf{|}\boldsymbol{Z}\mathbf{=C]}$$

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.

Association and Causation

- We are still suffering from the 'fundamental problem' and what we actually observe are conditional effects.

- The association is the different risk observed in two disjoint subsets of the population, defined by the actual exposure.

- Causation is the different risk in the entire population under the possible exposures.

The problem of confounding

- The difference Ave[$Y|Z$=T] − Ave[$Y|Z$=C] 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

$$\text{Ave}[Y(\text{C})|Z=\text{T}] = \text{Ave}[Y(\text{C})|Z=\text{C}]$$

$$\text{Ave}[Y(\text{T})|Z=\text{T}] = \text{Ave}[Y(\text{T})|Z=\text{C}]$$

- 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.

Treatment assignment mechanism

- If we have measurements on covariates, X, treatment assignment is said to be ignorable or exchangeable if $(Y(\mathrm{C}), Y(\mathrm{T})) \perp Z \mid X$ where "$\perp$" means statistically independent of".

- In words, treatment assignment (receipt) is ignorable if the two potential outcomes are jointly independent of Z given X.

- If we allow for the covariates, X, in an appropriate way then we can obtain unbiased (unconfounded) estimates of the treatment effects.

- So the key question is "why do people receive the treatment or exposure they receive?"

Does Association = Causation?

- In general, this does not hold.
- In randomised experiments, the groups are exchangeable since we assume an equal distribution of characteristics, and we assume there are no confounders.
- Then under randomisation, the associational measures are equal to the causal effect measures.
- This exchangeability does not hold when we have non-compliance, due to the non-random selection effects involved.

A ‘perfect’ randomised controlled trial

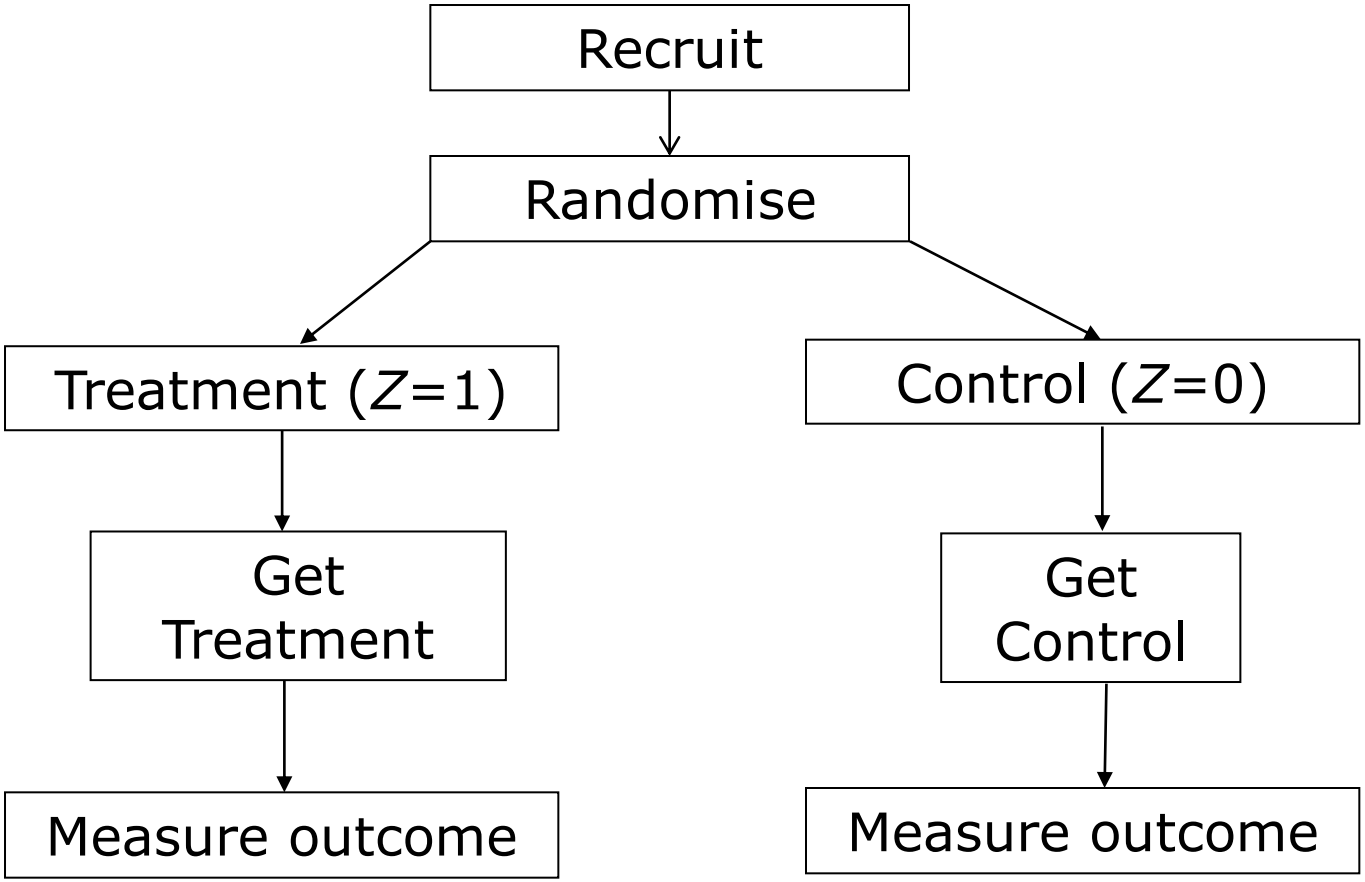

An ‘imperfect’ RCT

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.

A more realistic RCT

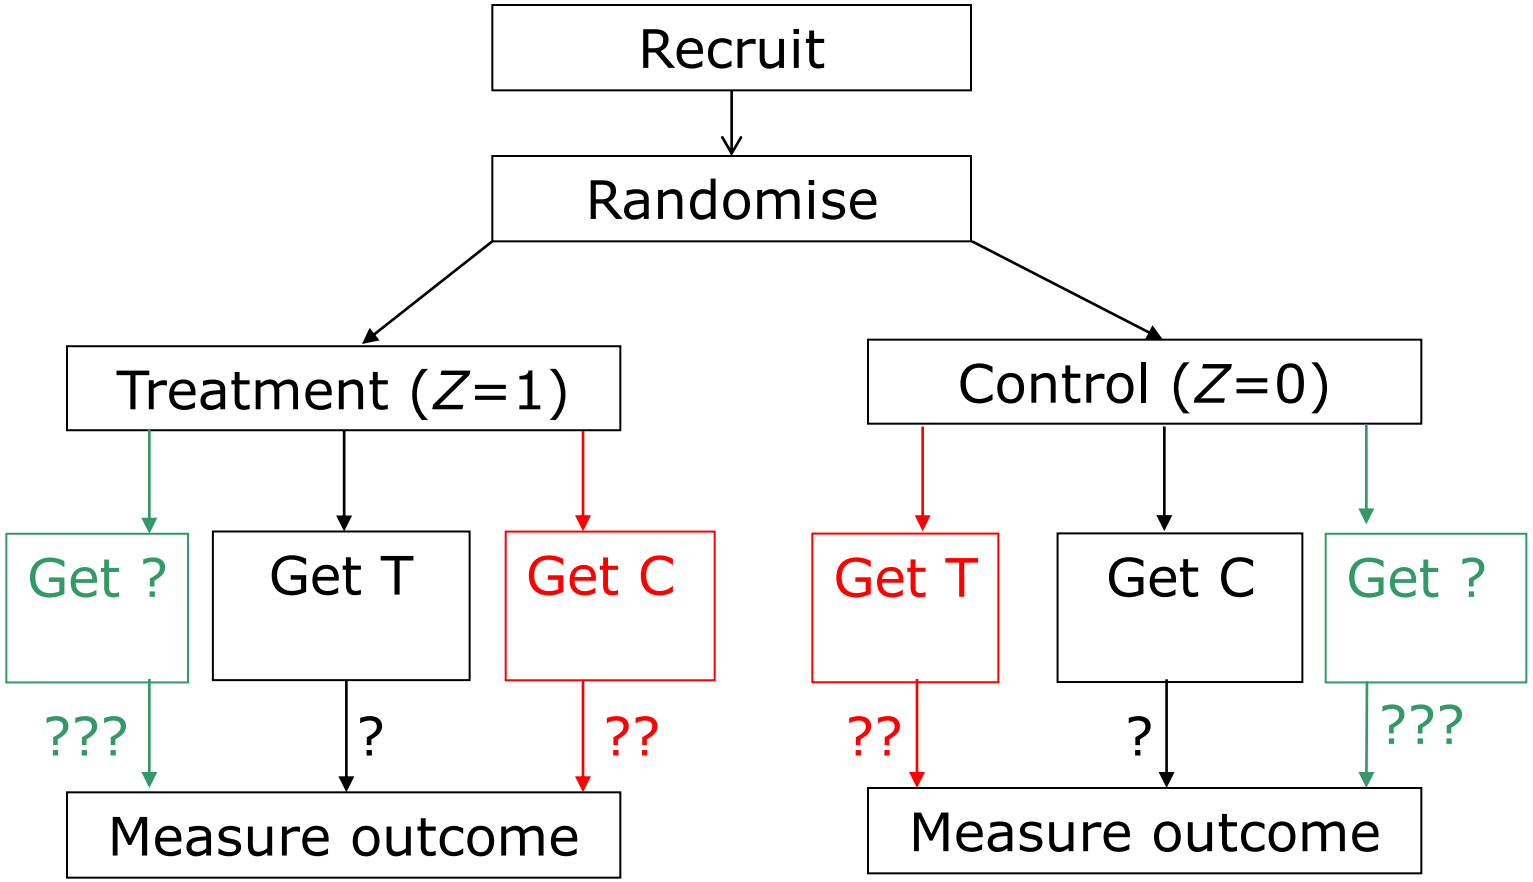

Switches

Changes to other or non-trial treatment

What is intention to treat (ITT)?

http://www.consort-statement.org/resources/glossary:

- "A strategy for analyzing data in which all participants are included in the group to which they were assigned, whether or not they completed the intervention given to the group.

- "Intention-to-treat analysis prevents bias caused by the loss of participants, which may disrupt the baseline equivalence established by random assignment and which may reflect non-adherence to the protocol."

- Compare subjects as randomised, regardless of what they actually received.

Problems in only focussing on ITT effects

- May be interested in these other questions, e.g. effectiveness of "being invited for screening" rather than the effect of being screened on those who are screened.
- Switching in particular makes ITT analysis conservative but for an equivalence trial makes ITT anti-conservative.
- Differences between trials (effect heterogeneity) may be the result of varying levels of compliance rather than heterogeneity in efficacy – problem for meta-analysis.
- Treatment by time interactions may be explained by increasing noncompliance/switching etc.
- Treatment-severity interactions may appear if switching associated with severity.

The Complier Average Causal Effect (CACE)

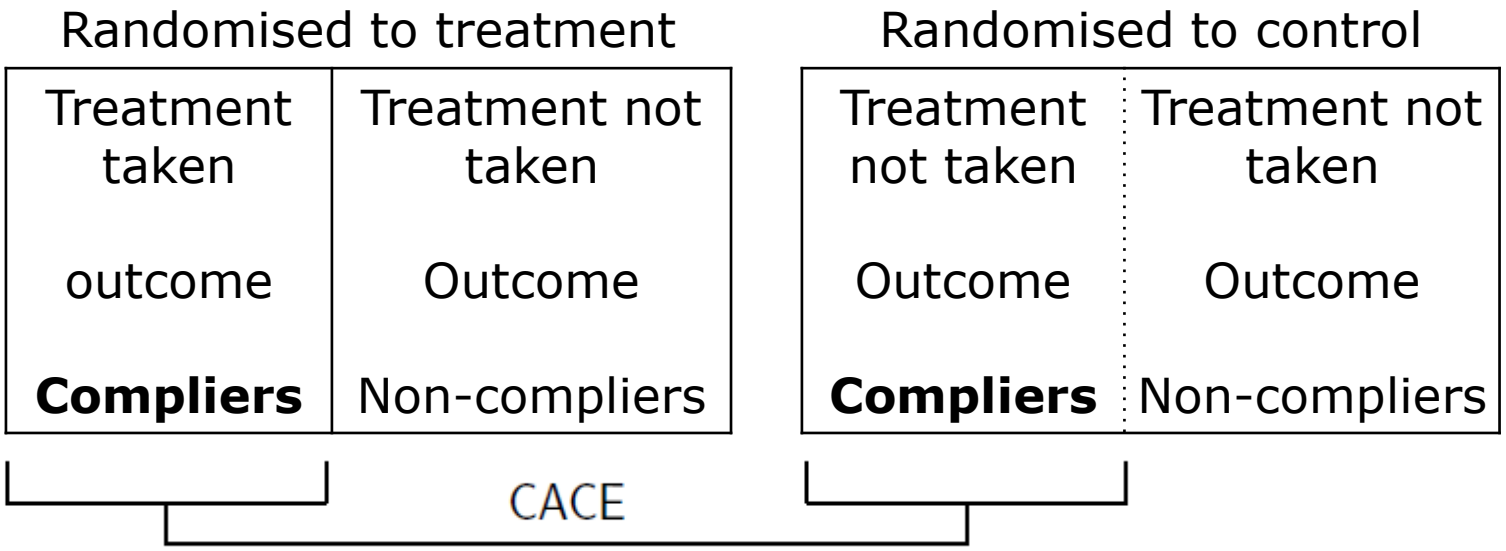

- The ***Complier-Average Causal Effect (CACE) estimate*** is the comparison of the average outcome of the compliers in the treatment arm with the average outcome of the comparable group of would-be compliers in the control arm.
- This is a randomisation-respecting estimate.
- It is the ITT effect in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

Simple mediation/mechanism diagram

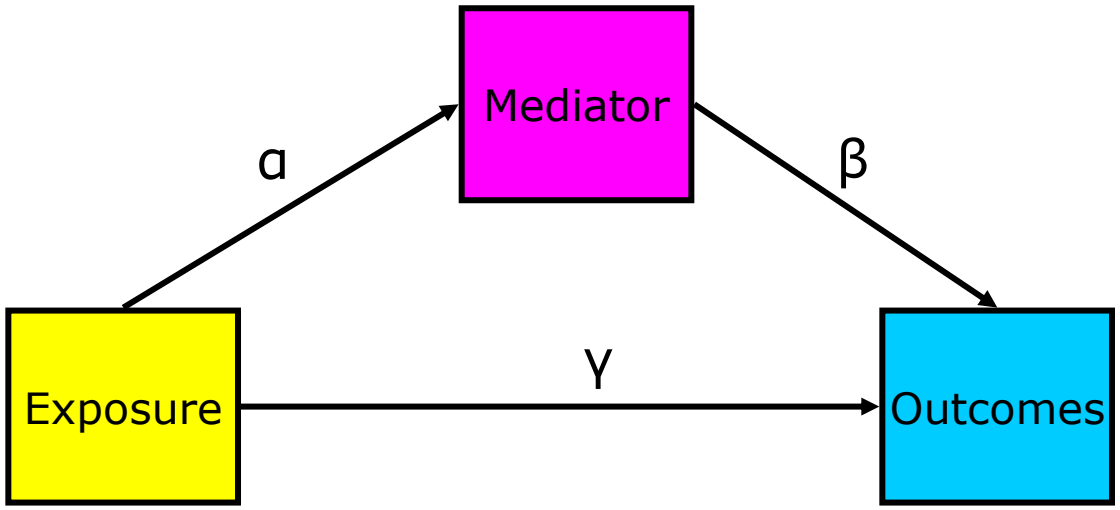

Total effect = direct effect + indirect effect

Mediation analysis and causal inference…

"Mediation analysis is a form of causal analysis…all too often persons conducting mediational analysis either do not realize that they are conducting causal analyses or they fail to justify the assumptions that they have made in their casual model."

David Kenny (2008), Reflections on Mediation, *Organizational Research Methods.*

Confounded mediation: estimating valid causal effects

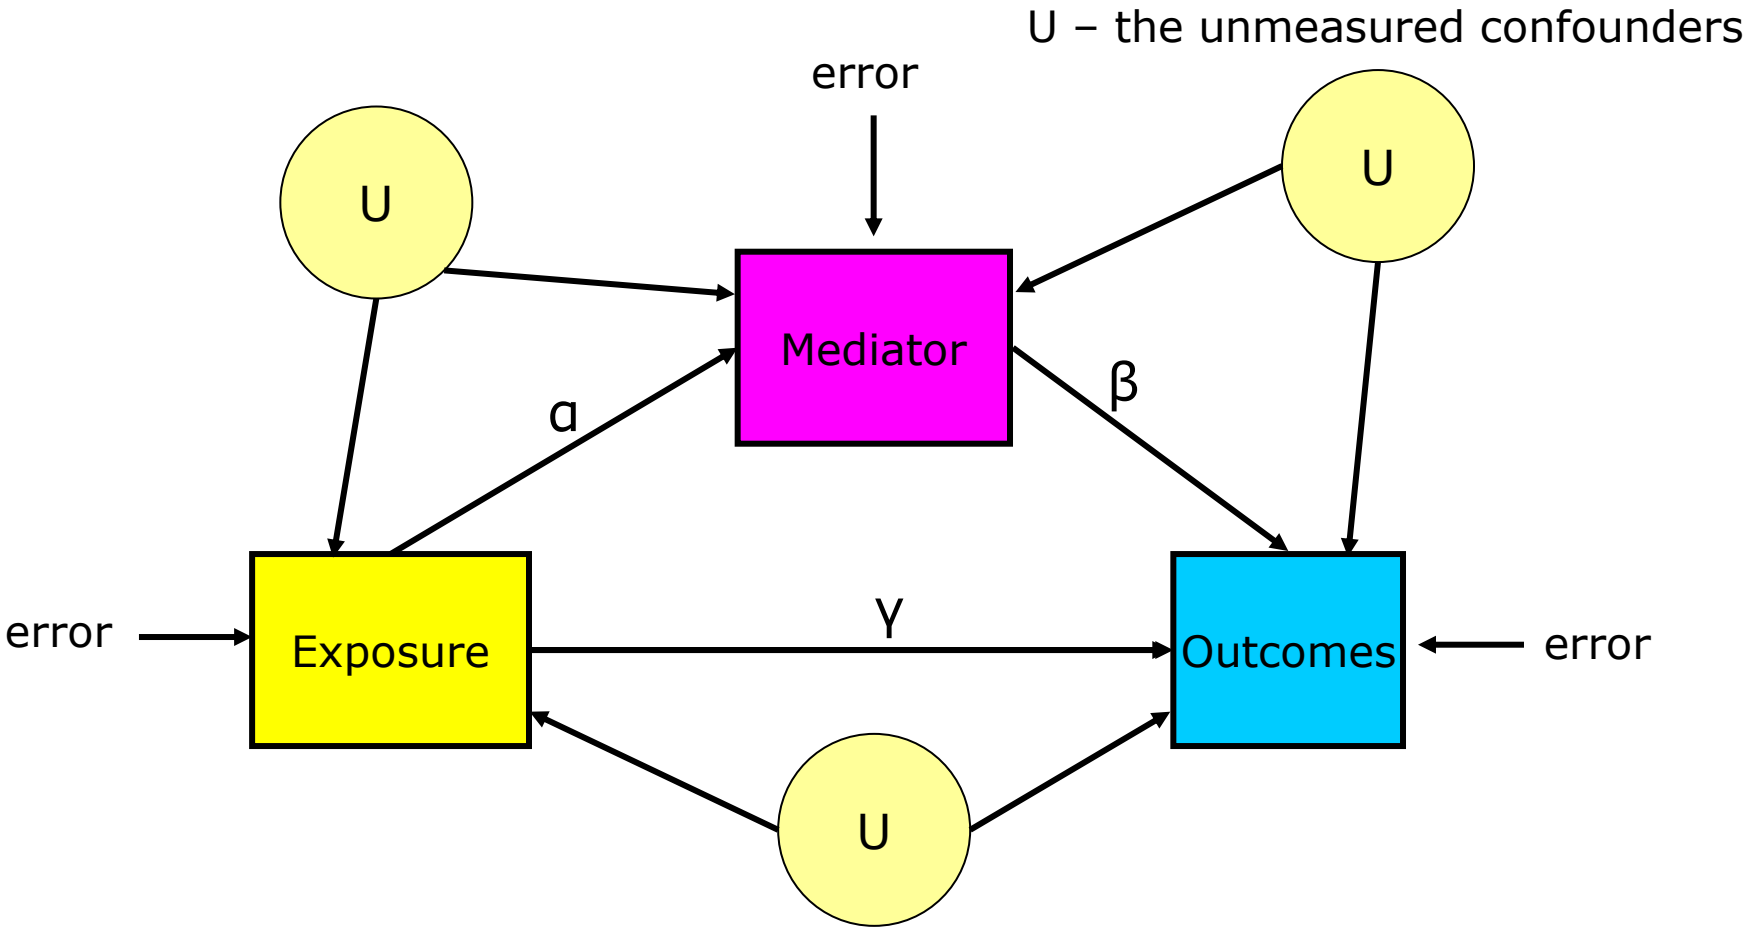

The basic underlying problem: estimating valid causal effects

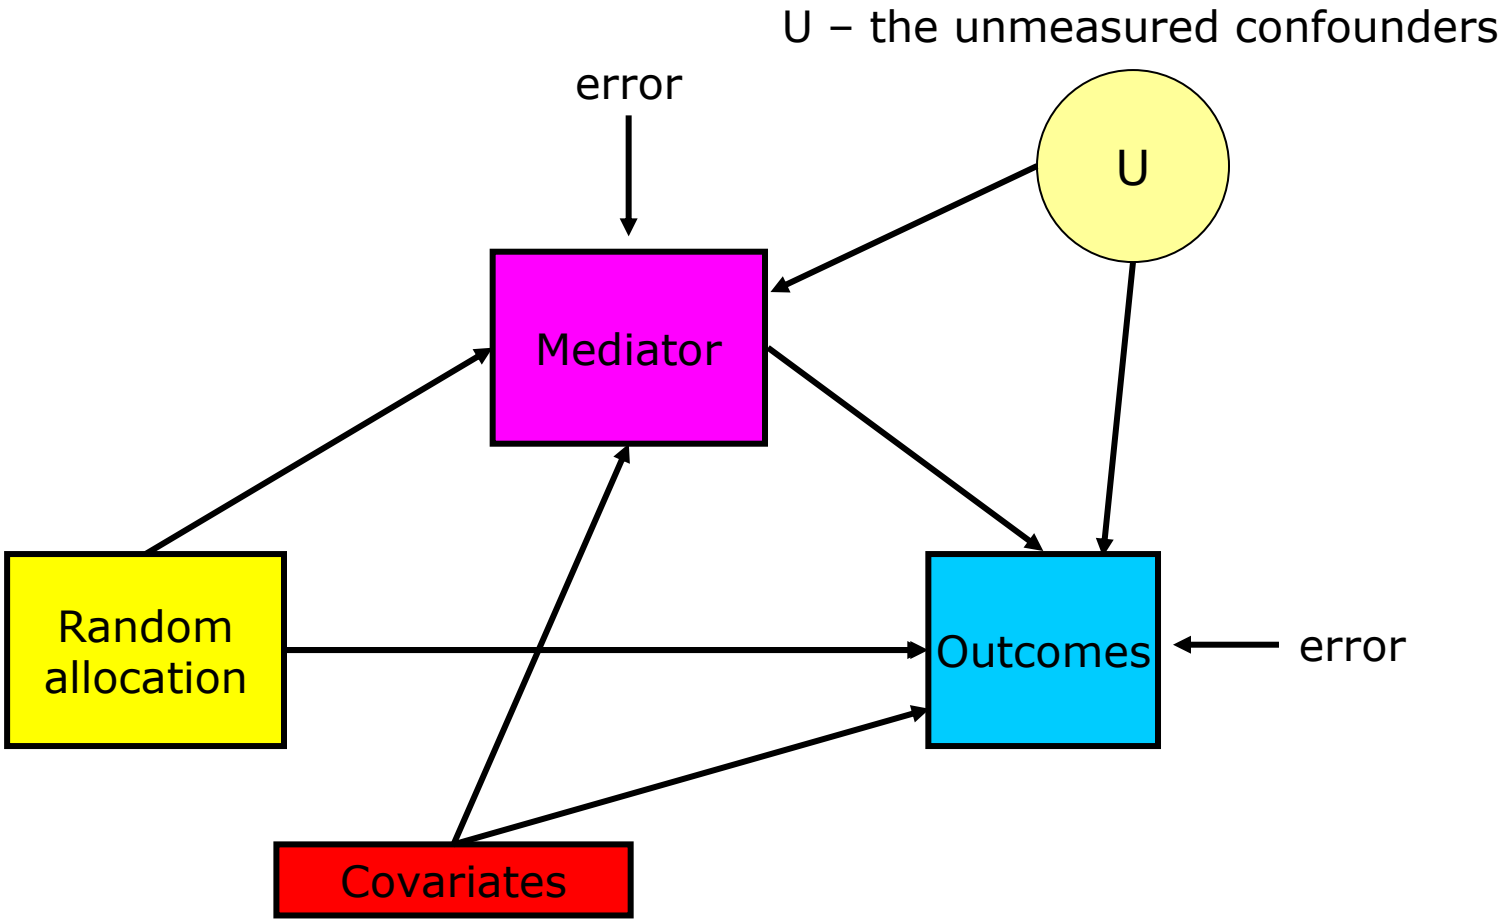

Statistical mediation analysis

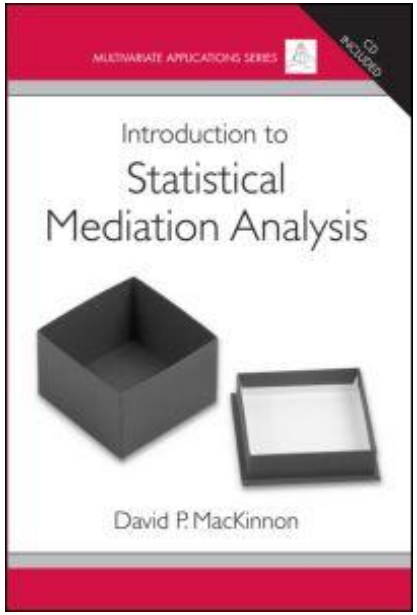

Dave MacKinnon

David Kenny

- Large literature on statistical mediation analysis, arising from Baron and Kenny (1986) and summarised by the recent monograph by David MacKinnon (2008).

- Extensive use of structural equation modelling including examples.

Statistical mediation analysis

- The Baron and Kenny procedure and subsequent estimation of the indirect effect can be appropriate, provided:
  - Continuous outcome and continuous mediator
  - All relevant confounders are included in all the models and there are no unmeasured confounders (e.g. excluding covariates)
  - Correct functional form (e.g. linearity)
  - There are no interactions between treatment and mediator on outcome.

- Using the bootstrap option is probably recommended for estimating the standard error of the indirect effect.

- This applies for the use of structural equation modelling more generally too.

Causal mediation analysis

- Statistical mediation (B&K) has three 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 the models are correctly specified.

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

- Large amount of methodological expertise on this topic at Manchester.

Causal mediation definitions: direct and indirect effects

- **(Pure) natural direct effect:** $Y_i(\mathrm{T},M_i(\mathrm{C}))-Y_i(\mathrm{C},M_i(\mathrm{C}))$
  - The direct effect of random allocation given $M(0)$, the 'natural' level of the mediator
- **(Total) natural indirect effect:** $Y_i(\mathrm{T},M_i(\mathrm{T}))-Y_i(\mathrm{T},M_i(\mathrm{C}))$
  - The effect of the change in mediator if randomised to receive treatment (i.e. Z=T).
- **Controlled direct effect:** $Y_i(\mathrm{T},\mathrm{m})-Y_i(\mathrm{C},\mathrm{m})$
  - Direct effect of randomisation on outcome at mediator level m.
- Total Effect = Natural direct effect + Natural indirect effect

(Pearl 2001;Robins & Greenland 1992).

A brief history of causal inference (3)

- Director of Program on Causal Inference at HSPH
- Extended the potential outcomes framework to longitudinal setting (repeated measures).
- This required a new methodology for estimating parameters using semi-parametric theory: the "G-family"
- Uses terminology 'counterfactuals' rather than potential outcomes.

Jamie Robins

A brief history of causal inference (4)

- Current and past members of the HSPH CI program.
- Developing estimation methods for lots of practical questions within the "G-Family"
  - G-estimation
  - G-formula
  - G-computation
- The most widely-cited and published group in CI.

From top left: Stijn Vansteelandt, Andrea Rotnitzky, Miguel Hernán, Tyler VanderWeele, Eric Tchetgen Tchetgen, Els Goetghebeur

Confounding adjustment

- **Standard Approach**: Model the probability of disease, taking into account the past exposure and past history of possible confounders using methods such as logistic or proportion hazards regression.

- These approaches may be biased when:

  1. There exists a time-dependent covariate that is a risk factor for, or predictor of, the event of interest and also predicts future exposure.

  2. Past exposure history predicts the subsequent level of covariate.

- These conditions will be true in many observational studies, and will always hold when there are time-dependent covariates that are simultaneously confounders and intermediate variables.

Jamie Robins (1986) – his first causal inference paper

Mathematical Modelling, Vol. 7, pp. 1393–1512, 1986
Printed in the U.S.A.

0270–0255/86 $3.00 + .00

A NEW APPROACH TO CAUSAL INFERENCE IN MORTALITY STUDIES WITH A SUSTAINED EXPOSURE PERIOD—APPLICATION TO CONTROL OF THE HEALTHY WORKER SURVIVOR EFFECT

JAMES ROBINS

Harvard School of Public Health
665 Huntington Avenue
Boston, MA 02115

(*Received* 19 *June* 1985; *revised* 12 *January* 1986)

286 WoS citations

Healthy Worker Survivor Effect

- "In occupational mortality studies date of termination of employment is both a determinant of future exposure (since terminated individuals receive no further exposure) and an independent risk factor for death (since disabled individuals tend to leave employment).

- When current risk factor status determines subsequent exposure and is determined by previous exposure, standard analyses that estimate age-specific mortality rates as a function of cumulative exposure may underestimate the true effect of exposure on mortality whether or not one adjusts for the risk factor in the analysis.

- This observation raises the question, which if any population parameters can be given **a causal interpretation** in observational mortality studies?"

- Abstract from Robins (1986)

Healthy Worker Survivor Effect

- This was first observed by Ogle (1885)
  - "the more vigorous occupations had relatively lower mortality rate as compared with the death-rates in occupations of an easier character or the unemployed".
- The term "healthy worker effect" was first used by McMichael et al (1974).
- Robins (1986) was the first to derive mathematically the concept of HWSE and propose a statistical solution.
  - Introduces concept of time-varying confounding

Time varying confounding

- Confounders could be fixed (usually at baseline) or time-varying.

- A variable is a time-dependent confounder if it predicts
  1. future treatment and
  2. future outcome, conditional on past treatment.

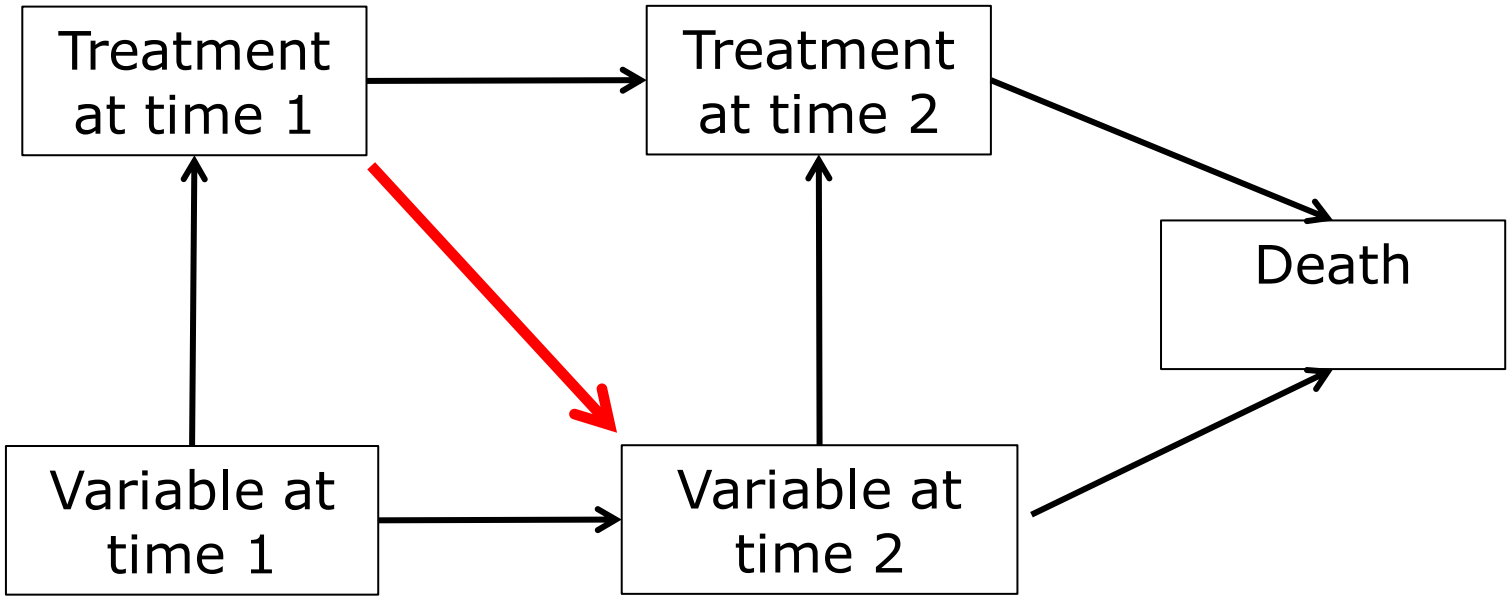

Classic example: LDL count in HIV

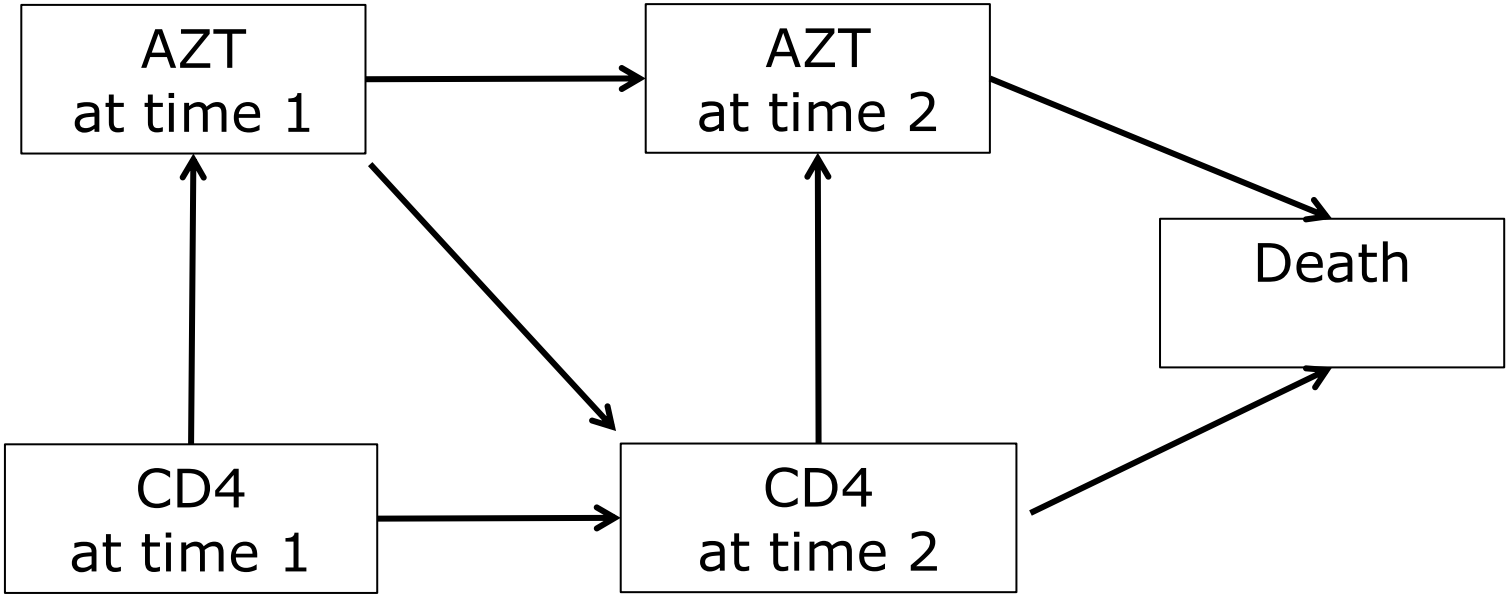

- Time-dependent confounding of CD4 lymphocyte count affected by previous treatment on risk of mortality from zidovudine (AZT).
- So CD4 count predicts HAART and HAART raises CD4 counts.

Example: Obesity and mortality

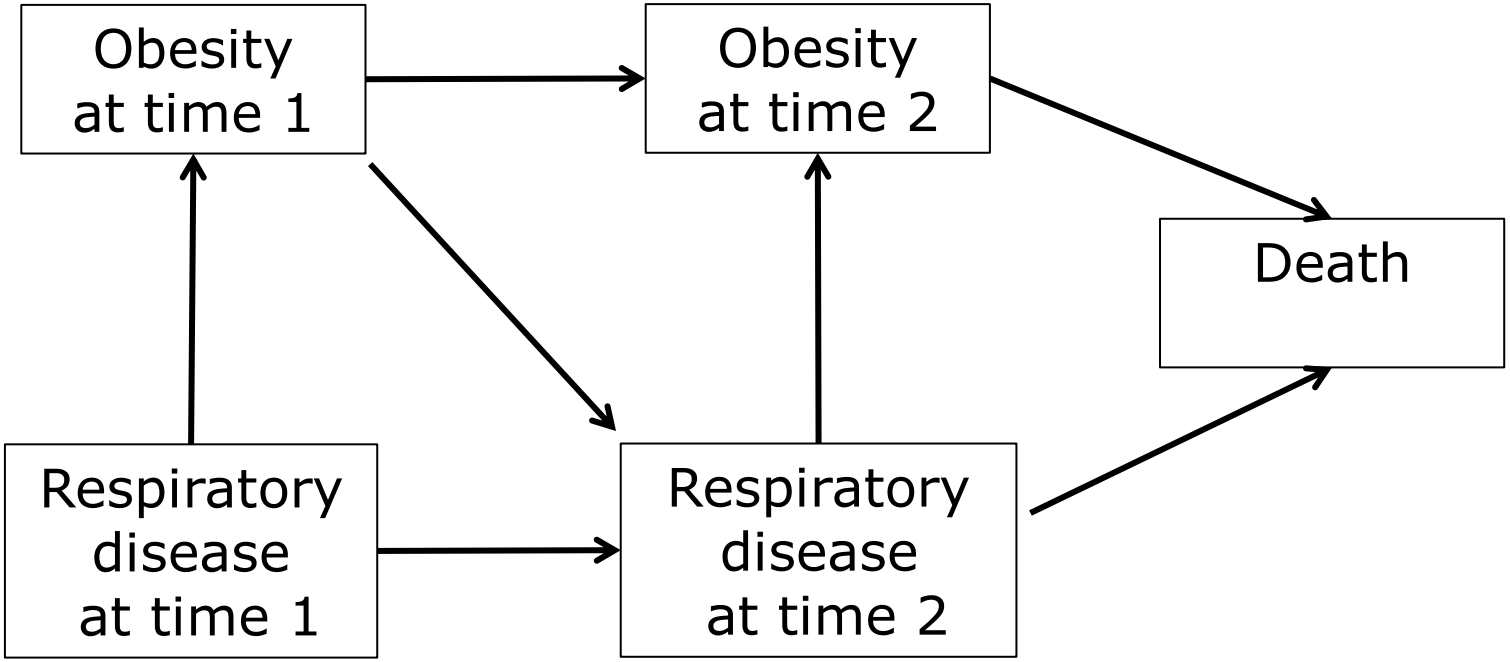

- Time-dependent confounding of respiratory disease affected by previous treatment on risk of mortality from some obesity measure.

Unmeasured confounding between variable and outcome

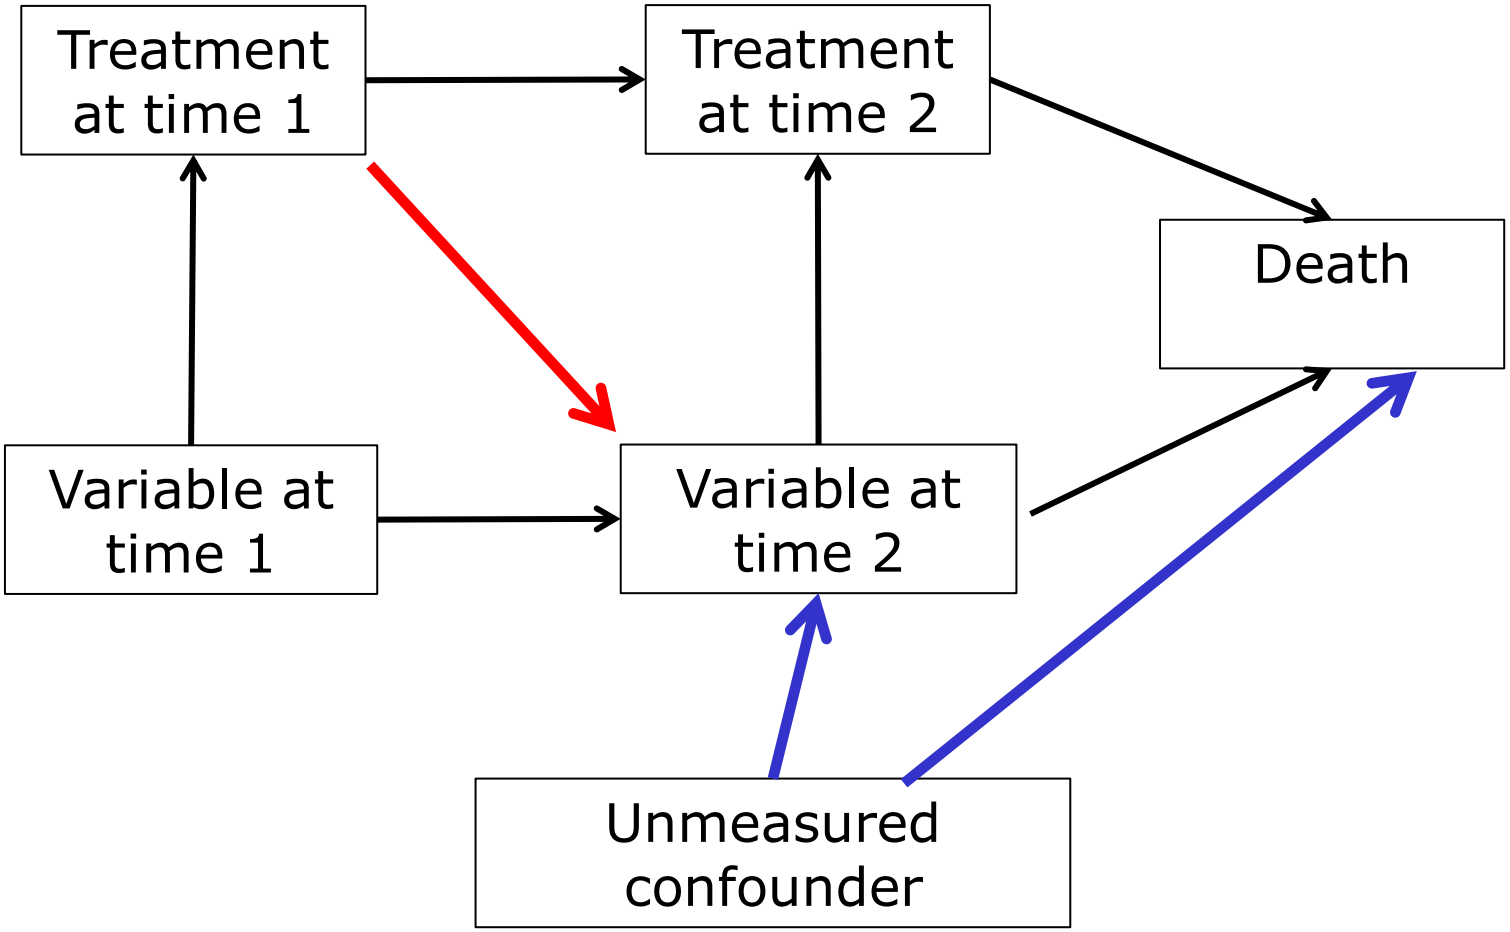

Unmeasured confounding between variable treatment and outcome

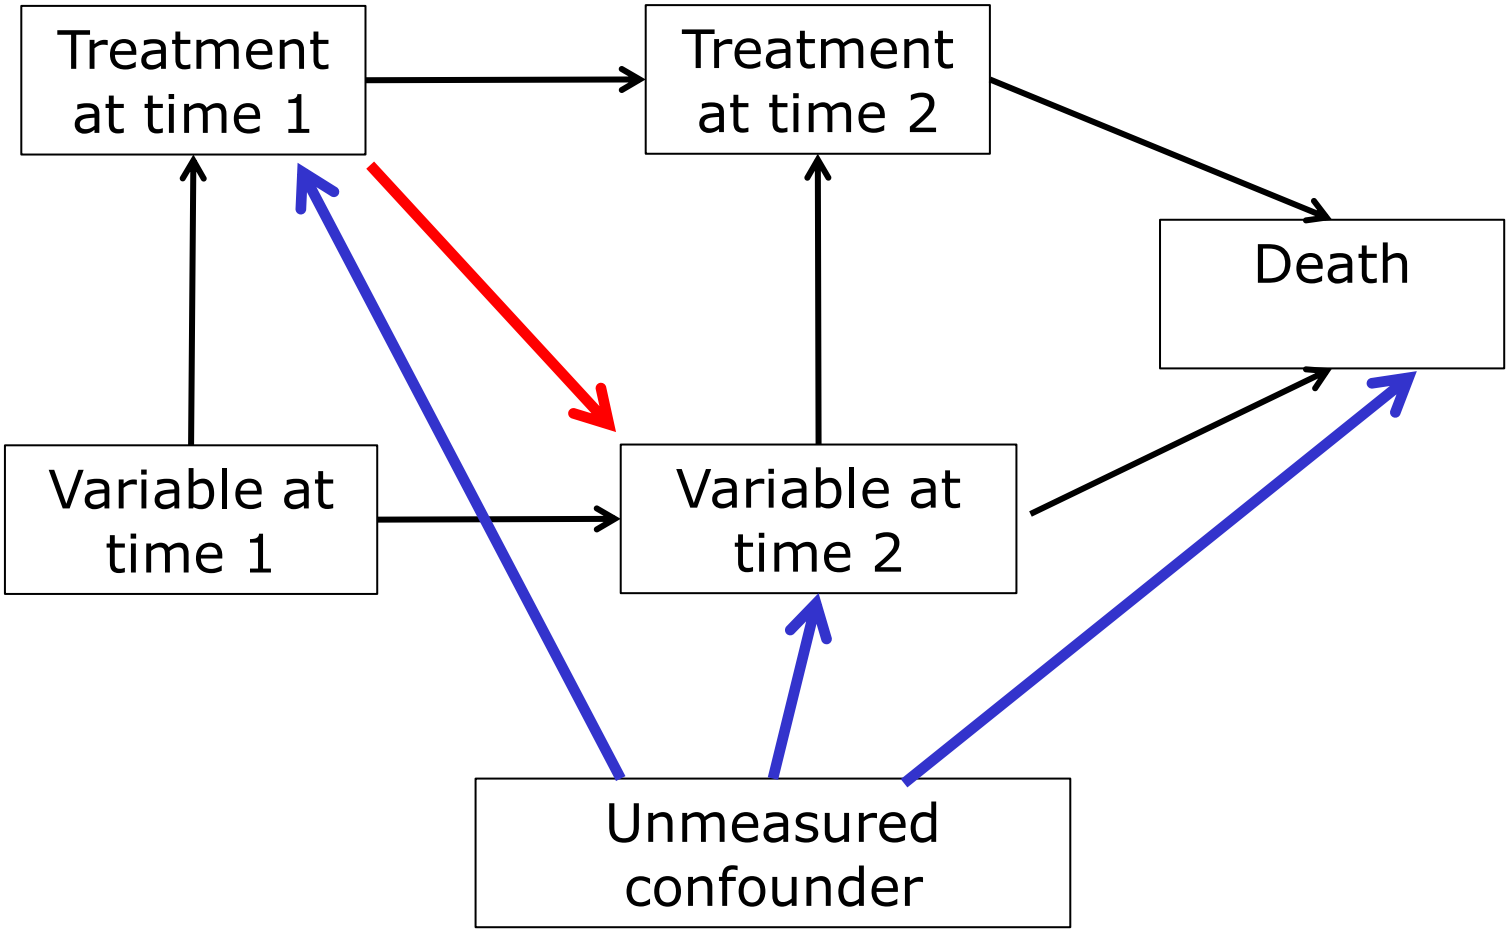

Controlling for a variable affected by treatment

- In standard regression methods, we would adjust for covariates by including them as covariates. This adjustment may fail when adjusting for confounding due to **measured confounders** and when the **treatment is time varying** because:

  1. The **confounder** may be a confounder for later treatment and should be adjusted for, but
  2. The **confounder** may be affected by earlier treatment, and so should not be adjusted for, according to the standard methods.

- The solution suggested is to adjust for time dependent covariates by using them to calculate **inverse probability weights** rather than including them as covariates.

Time varying confounding

- Confounders could be fixed (usually at baseline) or time-varying.

- A variable is a time-dependent confounder if it predicts
  1. future treatment and
  2. future outcome, conditional on past treatment.

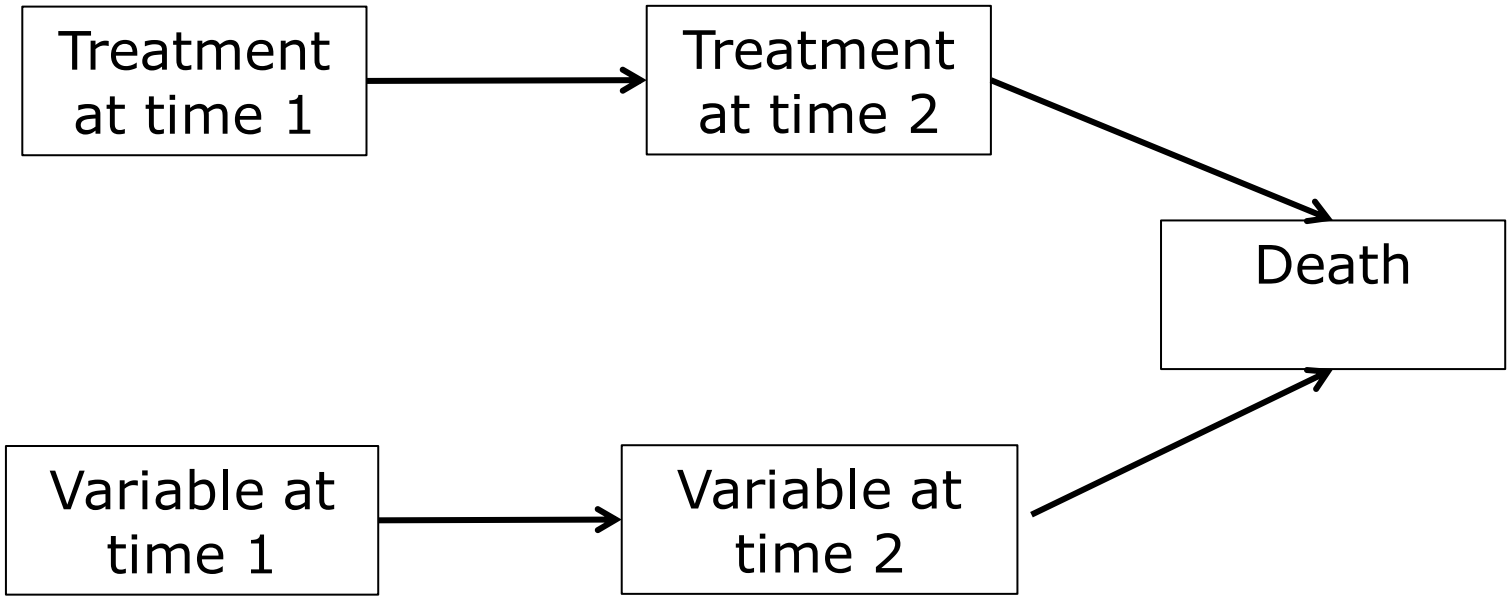

Marginal structural models: original article

Marginal Structural Models and Causal Inference in Epidemiology

James M. Robins,[1,2] *Miguel Ángel Hernán*,[1] *and Babette Brumback*[2]

In observational studies with exposures or treatments that vary over time, standard approaches for adjustment of confounding are biased when there exist time-dependent confounders that are also affected by previous treatment. This paper introduces marginal structural models, a new class of causal models that allow for improved adjustment of confounding in those situations. The parameters of a marginal structural model can be consistently estimated using a new class of estimators, the inverse-probability-of-treatment weighted estimators. (Epidemiology 2000;11:550–560)

Keywords: causality, counterfactuals, epidemiologic methods, longitudinal data, structural models, confounding, intermediate variables

Epidemiology: September 2000 - Volume 11 - Issue 5 - pp 550-560

715 WoS citations.

Marginal structural models: basic idea

- Suppose that at each timepoint t, we could create an identical copy of each patient.
- Then if the real patient received treatment, we would give the copy control and vice versa.
- We could then compare the patient to its copy.
- This solves confounding by matching: the patient is matched with the copy.
- Obviously this is impossible but we can use the idea to define the **counterfactuals** for each patient to be the outcomes for each of imaginary copies.

Marginal structural models: basic idea

- Idea: treat the counterfactuals as missing data and use inverse probability weighting.
- MSMs use inverse-probability weighting to deal with the unobserved ("missing") counterfactuals.
- We cannot adjust for confounders but using IPW, can re-weight the dataset so that treatment and covariates are unconfounded.

  i.e. that the mean covariate levels are the same between treated and untreated patients.
- So we can do a simple marginal analysis.
- Gives an average causal effect, but can be used to estimate average treatment effect on treated, average treatment effect on untreated.

Key assumption: Conditional Exchangeability

- In observational studies, if there exists a confounder X, we can stratify by this to create exchangeability within levels of X. This in effect creates a randomised experiment within the levels of X, i.e.

$$Y(Z) \perp Z \mid X=x \text{ for all } z$$

- However, this assumption cannot be tested.

- We can weight our population to produce this conditional exchangeability, thereby producing a randomised pseudo-population which allows us to use the association=causation argument from previously.

Where can MSMs be used?

- Comparative effectiveness research
- Mediation analysis
- Time-varying treatments
  - Non-compliance in randomised trials

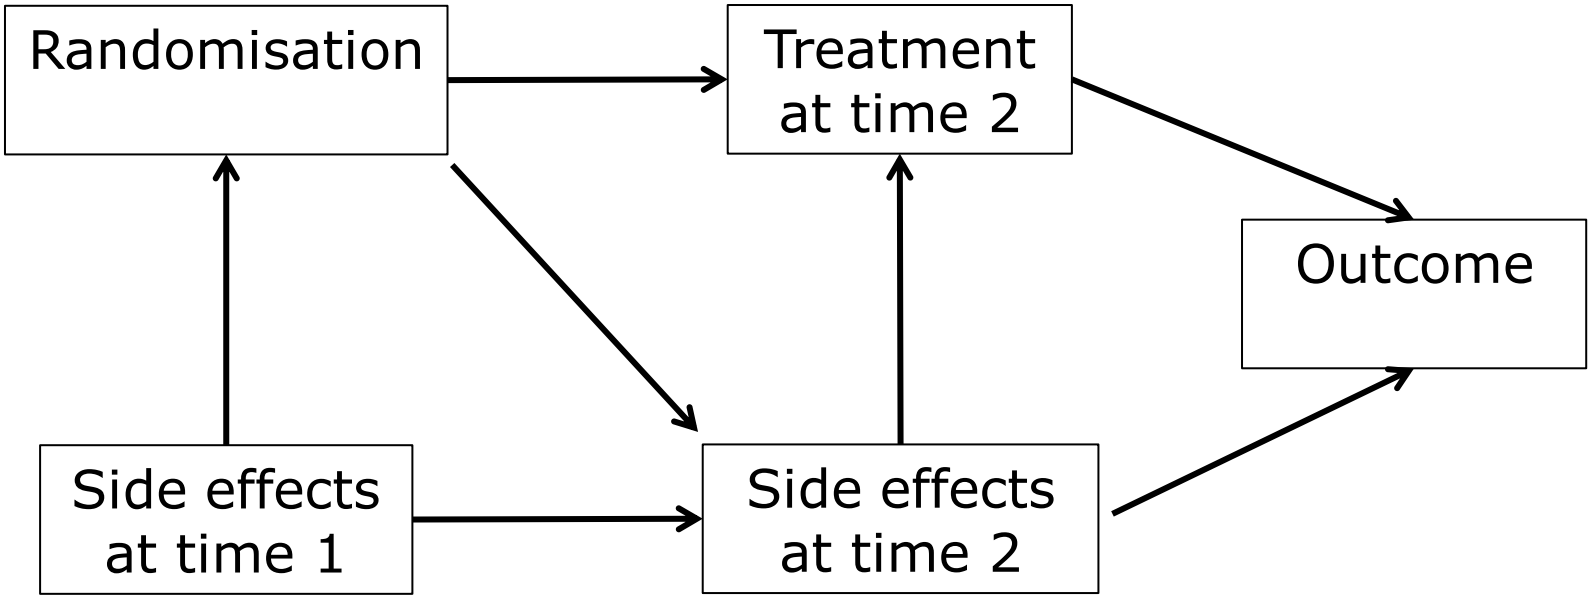

A brief history of causal inference (5)

- Developed a theory of causal and counterfactual inference based on graphical models and probabilistic reasoning.
- Derived a new method for determining relations between variables, known as ‘do-calculus’.
- Explores the link between counterfactuals and non-parametric structural equation models.

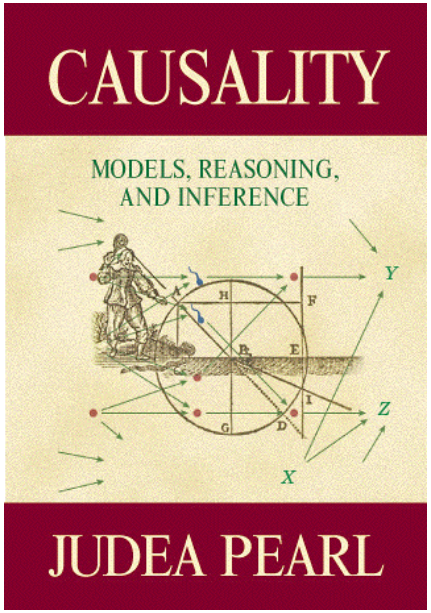

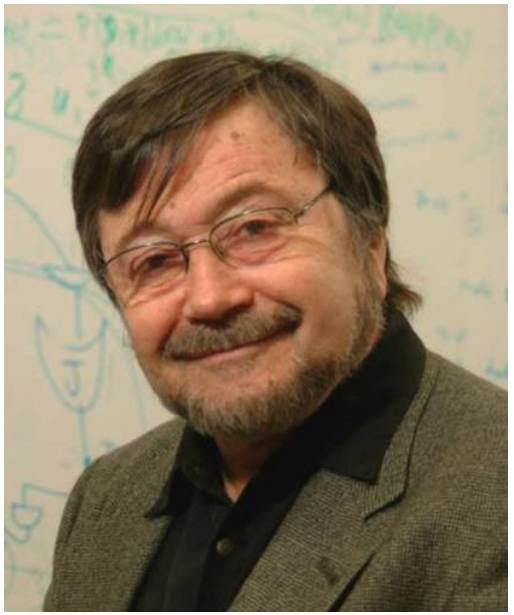

Judea Pearl

Path diagrams/Directed Acyclic Graphs

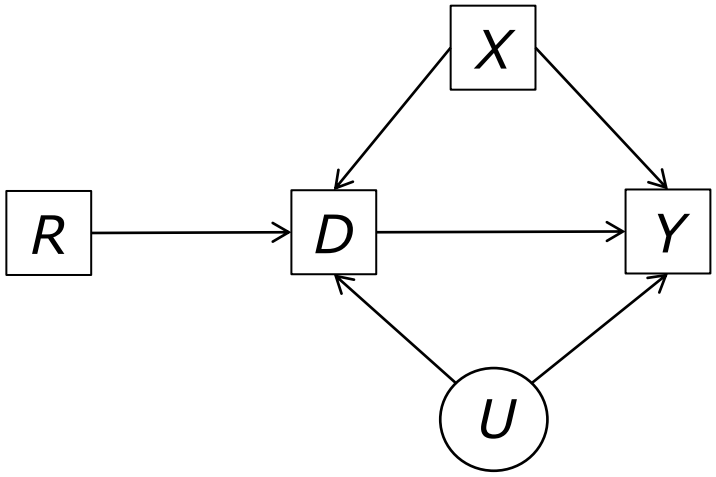

- Observed variables in squares, unobserved (latent) variables in circles.
- An arrow (directed link) between variables represents a causal effect.
- *X* is a measured confounder, *U* is an unmeasured confounder.
- Use the graph to read of conditional independencies:
  $Y \perp R \mid D,X,U$

Link with Pearl's do operator

$$\text{ITE} = \delta_i = Y_i|\text{do}(Z_i=1) - Y_i|\text{do}(Z_i=0)$$

where 'do(Z_i=1)' means that the Z_i is set to one by the investigator.

- This is mathematically equivalent to the description in terms of potential outcomes but has the advantage of emphasising the importance of experimental manipulation.
- "Doing" rather than just "Seeing".
- It emphasises the role of intervening on variables, i.e. removing all arrows into the variable in the DAG.

A brief history of causal inference (6)

- There is a group who argue against using the counterfactuals or potential outcomes framework.
- Dawid and colleagues propose for methods for causal inference without counterfactuals, mainly using decision theory, graphical models and stochastic modelling.

L-R: Carlo Berzuini, Phil Dawid, Vanessa Didelez

Objections to counterfactuals (Dawid, 2000)

- The concepts are deterministic rather than stochastic – the potential outcomes are fixed for each subject, and just waiting to be revealed according to the assignment mechanism.

- We can never learn anything about the joint distribution of $Y_i(1)$ and $Y_i(0)$ when they can never be observed together. In particular we can never learn anything about the covariance (correlation) of $Y_i(1)$ and $Y_i(0)$. Although we can in principal estimate the average of the individual treatment effects, we cannot estimate their variance.

Is the terminology important?

"Personally I see the different formalisms as different 'languages'. The French language may be best for making love whereas the Italian may be suitable for singing, but both are indeed possible…"

Lauritzen: Scandinavian Journal of Statistics 2004 Vol. 31 p189

Some recent volumes on causal inference

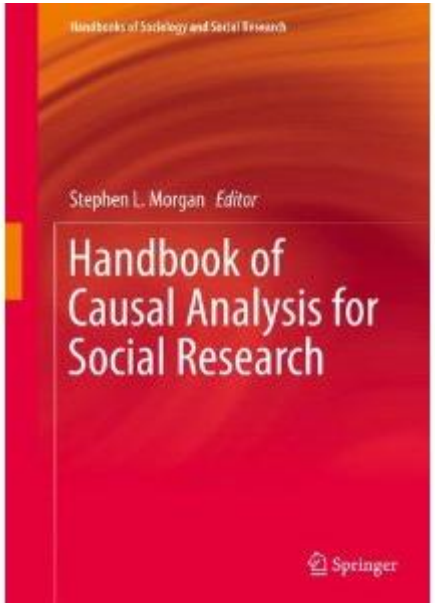

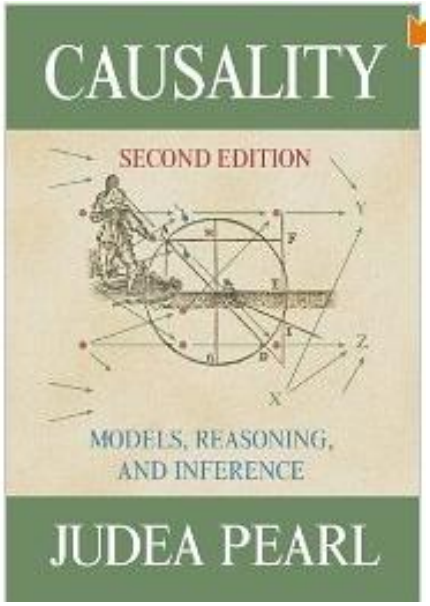

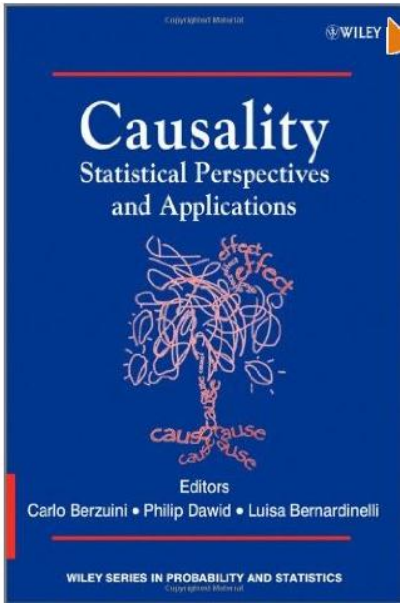

- Plus forthcoming volumes:
  - 'Causal Inference' by Robins and Hernán.
  - 'Causal Inference in Statistics and Social Sciences' by Imbens and Rubin.

New Journal of Causal Inference

- Journal of Causal Inference (JCI) publishes papers on theoretical and applied causal research across the range of academic disciplines that use quantitative tools to study causality.

- Journal of Causal Inference aims to provide a common venue for researchers working on causal inference in biostatistics and epidemiology, economics, political science and public policy, cognitive science and formal logic, and any field that aims to understand causality.

Expertise at the Centre for Biostatistics

http://www.population-health.manchester.ac.uk/biostatistics

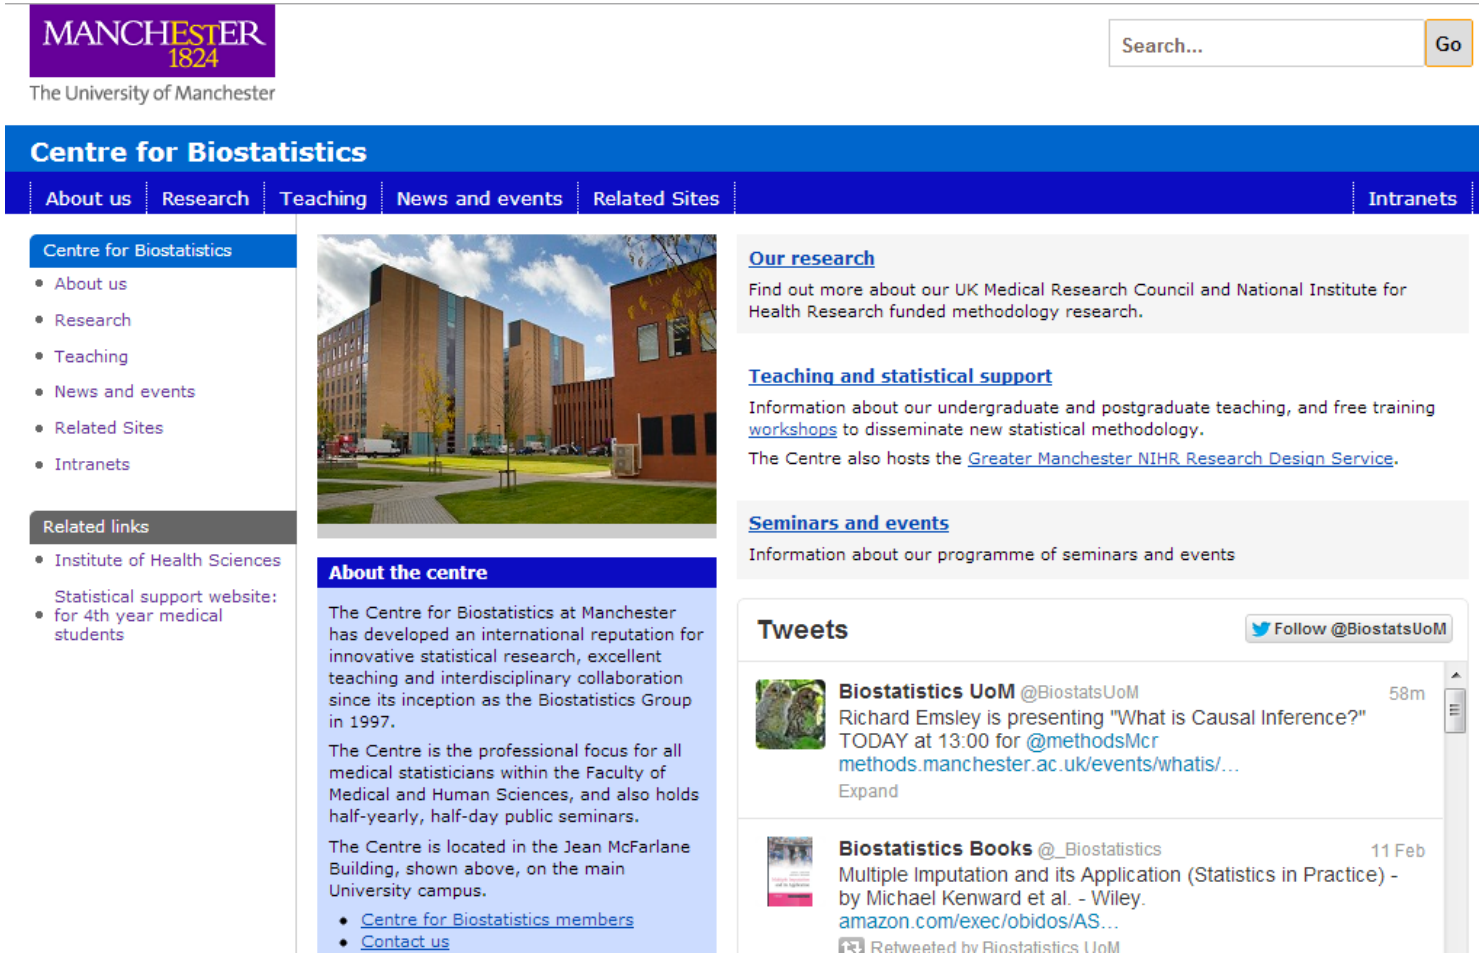

UK-Causal Inference Meeting (UK-CIM)

- https://sites.google.com/site/ukcausalinferencemeeting/
- First UK wide meeting on causal inference, organised by the Centre for Biostatistics.
- Tuesday 14th May - Wednesday 15th May 2013 at The University of Manchester.
- Theme: “Causal Inference in Health and Social Sciences”.
- Stijn Vansteelandt confirmed as Keynote Speaker
- Registration now open: £60 academics, £30 students.