University of BRISTOL

UK Causal Inference Meeting 2015

Conference Programme: 15th-17th April 2015

MRC Integrative Epidemiology Unit

UK-CIM

Organiser:
Nayia Constantinou (Bristol)
UK-CIM Committee

Administrative support:
Liz Wainwright (Bristol)

Welcome to the UK Causal Inference Meeting (UK-CIM)

This is the third meeting in an initiative to organise a regular UK based meeting on causal inference as a collaborative effort across the methodology research community in the health and social sciences. The first meeting was held in Manchester in May 2013 and the second meeting was held in Cambridge in April 2014.

UK-CIM aims to: provide a forum for people interested in causal inference to meet informally; for early career researchers to highlight their work; and to offer opportunities for networking to foster future research opportunities and collaborations.

The theme of the meeting is "Causal Inference in Health, Economic and Social Sciences". Causal inference is broadly defined, and the focus is on methodology and challenging applications, though presentations relating to interesting applications that highlight necessary methodological extensions are also encouraged

UK-CIM steering group:

Richard Emsley (Chair), University of Manchester

Mike Brewer, University of Essex

Stephan Burgess, University of Cambridge

Paul Clarke, University of Essex

Nayia Constantinou, University of Bristol

Rhian Daniel, London School of Hygiene and Tropical Medicine

Bianca De Stavola, London School of Hygiene and Tropical Medicine

Vanessa Didelez, University of Bristol

Sara Geneletti, London School of Economics

Richard Grieve, London School of Hygiene and Tropical Medicine

Andrew Pickles, Institute of Psychiatry, Psychology and Neuroscience, KCL

Nigel Rice, University of York

Nuala Sheehan, University of Leicester

Kate Tilling, University of Bristol

Ian White, MRC Biostatistics Unit, Cambridge

SuSTaIN

SuSTaIn is Statistics underpinning Science, Technology and Industry, a programme with the ambitious goal of strengthening the discipline of Statistics in the UK, by equipping it to face the challenges of future applications. Thus the focus is on rigorous and innovative new theory and methodology – core statistics for the 21st century – aimed at and stimulated by generic challenges raised by the 'data revolution', in areas as diverse as genomics, astronomy, telecommunications and finance.

It is funded principally by a 3.5million Science and Innovation award from EPSRC, and partly by the University of Bristol, and runs from 2006 to 2016.

MRC Integrative Epidemiology Unit at the University of Bristol

At this year's UK-CIM, the MRC Integrative Epidemiology Unit at the University of Bristol (MRC IEU) funds one invited speaker, Prof Eric Tchetgen Tchtegen (Harvard). The MRC IEU conducts advanced population health science research to improve our understanding of the biological and environmental factors that can underlie and trigger common disease.

Wednesday 15th April 2015
Venue: Lecture theatre 2, Chemistry Building

09.00-10.00 Arrival and registration (refreshments)

10.00-10.10 UK-CIM Introduction by Richard Emsley

10.10-10.20 Welcome talk by Nayia Constantinou

10.20-10.30 SuSTaIn information talk by Guy Nason

10.30-12.00 Session 1 (Instrumental Variables)
Chair: Stephen Burgess, University of Cambridge

10.30-10.50 *Jack Bowden,* MRC Biostatistics Unit, Cambridge
Mendelian randomization with invalid instruments: bias from pleiotropy detected by simple graphical and statistical tests.

11.50-11.10 *Hyunseung Kang,* University of Pennsylvania
Robust Confidence Intervals with invalid instruments.

11.10-11.30 *Cedric E. Ginestet*, King's College London
Convex combination of ordinary least squares and two-stage least squares estimators.

11.30-12.00 Invited talk 1: *Eric Tchetgen Tchetgen,* Harvard University
Unification of the instrumental variable approach for causal inference and missing data.

12.00-13.30 Lunch break

13.30-15.00 Session 2 (Instrumental Variables and Causal Search)
Chair: Nuala Sheehan, University of Leicester

13.30-13.50 *Karla Diaz-Ordaz,* London School of Hygiene and Tropical Medicine
Instrumental variable approaches for estimating causal effects in settings with multivariate outcomes.

13.50-14.10 *Ditte Nørbo Sørensen,* University of Copenhagen
A structural Cox model for the causal effect of exposure among the exposed in an IV setting.

14.10-14.30 *Ricardo Silva,* University College London
Relaxing the assumptions of causal discovery algorithms.

14.30-15.00 Invited talk 2: *Elizabeth L. Ogburn*, Johns Hopkins University

Causal and statistical inference with social network data: Massive challenges and meager progress.

15.00-15.30 Coffee break

15.30-16.30 Session 3
Chair: Rhian Daniel, London School of Hygiene and Tropical Medicine

15.30-16.30 Keynote talk 1: *Philip Dawid* , University of Cambridge
From Statistical Evidence to Evidence of Causality

16.30-18.00 Poster presentation & wine reception

Thursday 16th April 2015
Venue: Lecture theatre 2, Chemistry Building

9.30-11.00 Session 4 (Survival and Mediation)

Chair: Kate Tilling, University of Bristol

09.30-09.50 *Odd O. Aalen*, University of Oslo
Treatment effect of the treated: Understanding time-dependent confounding in survival analysis.

09.50-10.10 *Susanne Strohmaier*, University of Oslo
A simple to implement algorithm for natural direct and indirect effects in survival studies with a repeatedly measured mediator.

10.10-10.30 *Sjouke Vandenberghe*, Ghent University

Mediation analysis of randomised experiments.

10.30-11.00 Invited talk 3: *Arvid Sjölander,* University Stockholm
Bounds on biological/causal interactions

11.00-11.30 Coffee break

11.30-12.30 Session 5

Chair: Sara Geneletti, London School of Economics and Political Science

11.30-12.30 Keynote talk 2: *Kosuke Imai,* Princeton University
Causal Interaction in High Dimension

12.30-14.00 Lunch break

14.00-13.30 Session 6 (Longitudinal and Double Robustness)

Chair: Vanessa Didelez, University of Bristol

14.00-14.20 *Karel Vermeulen,* Ghent University

Bias-Reduced Doubly Robust Estimation.

14.20-14.40 *Michael P Wallace,* McGill University
Doubly-robust dynamic treatment regimen estimation via weighted least squares.

14.40-15.00 *Marco Doretti,* University of Perugia
Tackling non-ignorable dropout in the presence of time-varying confounding.

15.00-15.30 Invited talk 4: *Erica E.M. Moodie*, McGill University
Estimating the optimal treatment sequence for graft-versus-host-disease following bone marrow transplantation

15.30-16.00 Coffee break

16.00-17.30 Session 7 (Survival, Health and Work)

Chair: Richard Emsley, University of Manchester

16.00-16.20 *Kjetil Røysland,* University of Oslo
Causal inference in survival analysis: An example from prostate cancer.

16.20-16.40 *Andrew M. Jones*, University of York
Acute health shocks and labour market exits.

16.40-17.00 *Jon Michael Gran,* University of Oslo
Causal inference in a multi-state model for sickness absence and return to work.

17.00-17.30 Invited talk 5: *José R. Zubizarreta,* Columbia University
Stable Weights that Balance Covariates for Causal Inference and Estimation with Incomplete Outcome Data

19.00 Conference dinner at the Bristol museum

Friday 17th April 2015
Venue: Lecture Theatre 1, Chemistry Building

9.30-11.00 Session 8 (Regression Discontinuity Designs)
Chair: Paul Clarke, University of Essex

09.30-09.50 *Nayia Constantinou*, University of Bristol
A formal treatment of Regression Discontinuity Designs.

09.50-10.10 *Sara Geneletti,* London School of Economics and Political Science
Regression discontinuity designs: The challenge of binary outcomes.

10.10-10.30: *Aidan G. O'Keeffe*, University College London
Dynamic Causal Inference for a binary outcome in a Regression discontinuity design using local independence.

10.30-11.00 Invited talk 6: *Victor Chernozhukov,* MIT
Program Evaluation with High-Dimensional Data

11.00-11.30 Coffee break

11.30-12.45 Session 9
Chair: Ian White, University of Cambridge

11.30-12.30 Keynote talk 3: *Els Goetghebeur*, Ghent University
Instrumental variables and survival analysis

12.30–12.45 Final reflections

12.45-14.00 Lunch break

End of the meeting

ORAL PRESENTATIONS
(in order of appearance)

***Mendelian randomization with invalid instruments: bias from pleiotropy detected by* simple graphical and statistical tests**

Jack Bowden, George Davey Smith** and Stephen Burgess****
**MRC Biostatistics Unit, Cambridge.*
***MRC Integrative Epidemiology Unit, University of Bristol.*
****Department of Public Health and Primary Care, University of Cambridge.*

Background: The number of Mendelian randomization analyses including large numbers of genetic variants is rapidly increasing. This is due to the proliferation of genome-wide association studies, and the desire to obtain more precise estimates of causal effects. However, some genetic variants may not be valid instrumental variables, in particular due to pleiotropy.
Methods: We view Mendelian randomization with multiple instruments as a meta-analysis, and show that bias caused by pleiotropy can be regarded as analogous to publication bias. Causal estimates using each instrument individually can be displayed visually by a funnel plot to assess potential asymmetry. Egger's test, a tool to detect publication bias, can be used to test for bias from pleiotropy. Under the assumption that the association of each genetic variant with the exposure is independent of the direct effect of the variant on the outcome (not via the exposure), Egger's test gives a valid test of the null causal hypothesis and a consistent causal effect estimate even when *all* the genetic variants are invalid instrumental variables.
Results: We illustrate the use Egger's test by re-analysing two recently published Mendelian randomization studies of the causal effect of height on lung function, and the causal effect of blood pressure on coronary artery disease risk.
Conclusions: Egger's test can detect some violations of the standard instrumental variable assumptions, and test for causal inference when these assumptions are violated. This provides a sensitivity analysis for the robustness of the findings from a Mendelian randomization investigation.

Robust Confidence Intervals with invalid instruments

Hyunseung Kang
Department of Statistics, University of Pennsylvania

Instrumental variables have been widely used to estimate the causal effect of a treatment on an outcome in the presence of unmeasured confounding. Existing confidence intervals for causal effects based on instrumental variables assumes that all of the putative instrumental variables are valid; a valid instrumental variable is a variable that affects the outcome only by affecting the treatment and is not related to unmeasured confounders. However, in practice, some of the putative instrumental variables are likely to be invalid. The paper presents a simple and general approach to construct a robust confidence interval that is robust to possibly invalid instruments. The robust confidence interval has theoretical guarantees on having the correct coverage. The paper also shows that the robust confidence interval outperforms traditional confidence intervals popular in instrumental variables literature when invalid

instruments are present. The new approach is applied to a study of the causal effect of income on food expenditures.

Convex combination of ordinary least squares and two-stage least squares estimators

Cedric E. Ginestet (1), Richard Emsley (2), and Sabine Landau (1)
(1) Biostatistics Department, Institute of Psychiatry, Psychology and Neuroscience, King's College London.
(2) Centre for Biostatistics, Institute of Population Health, University of Manchester.

In the presence of confounders, the ordinary least squares (OLS) estimator is known to be biased. This problem can be remedied by using the two-stage least squares (TSLS) estimator, based on the availability of valid instrumental variables (IVs). However, this reduction in bias is offset by an increase in variance. Under standard assumptions, the OLS has a larger bias than the TSLS estimator. Moreover, one can prove that the sample variance of the OLS estimator is no greater than the one of the TSLS.
Therefore, it is natural to ask whether one could combine the desirable properties of the OLS and TSLS estimators. Such a trade-off can be achieved through a convex combination of these two estimators, thereby producing our proposed convex least squares (CLS) estimator. The relative contributions of the OLS and TSLS estimators are here chosen to minimize the mean squared error (MSE) of their convex combination. This proportion parameter is proved to be unique, whenever the OLS and TSLS differ in MSEs. In particular, we also show that this proportion parameter can be estimated from the data, and that the resulting CLS estimator is consistent.

The finite-sample properties of this estimator are investigated using Monte Carlo simulations, in which we independently vary the amount of confounding and the strength of the instrument. Overall, the CLS estimator is found to outperform the TSLS estimator in terms of MSE. The method is also applied to a real-world data set from econometrics.

Unification of the instrumental variable approach for causal inference and missing data

Eric Tchetgen Tchetgen
Harvard University

Unobserved confounding is a well-known threat to causal inference with observational data. Likewise, selection bias can arise in the presence of missing data if there is an unobserved common cause of the missingness process and the outcome subject to missingness. An instrumental variable for unobserved confounding (IV-C) is a pre-exposure correlate of exposure known to only affect the outcome through its association with exposure. Likewise an instrumental variable for missing data (IV-M) is a predictor of missingness which is otherwise independent of the outcome in the underlying population. We give a general necessary and sufficient condition for nonparametric identification with an IV in settings (IV-C) or (IV-M), thus providing a unification of identification for causal inference and missing data with an IV. The approach equally applies for discrete or continuous IV and outcome. Interestingly, the proposed approach provides an elegant solution to the identification problem of the marginal effect of treatment on the treated with an IV which has been a longstanding problem in causal inference. For statistical inference incorporating high dimensional

covariates, we present generalizations of inverse-probability weighting, outcome regression and doubly robust estimation with an instrumental variable that equally apply to IV-C and IV-M. We illustrate the approach with several empirical examples.

Instrumental variable approaches for estimating causal effects in settings with multivariate outcomes

Karla Diaz-Ordaz (presenting author), Angelo Franchini, Richard Grieve
London School of Hygiene and Tropical Medicine

In randomised controlled trials that have non-compliance with the treatment assigned, policy makers require unbiased estimates of the causal effect of the treatment received. Instrumental variable (IV) approaches provide complier average causal effects (CACE) estimates. Common IV methods such as two-stage least squares (2SLS) have not been extended to settings with multivariate outcomes.
We propose a three-stage least squares (3SLS) regression approach, whereby estimates from the first stage regression of treatment received conditional on assignment, feed into a seemingly unrelated regression (SUR) system of equations that recognise the correlation between the outcomes. We also develop Bayesian IV approaches which jointly model the effects of random assignment on treatment received, and the bivariate outcome, which here is assume to be cost-effectiveness. We also apply 2SLS individually to each outcome, for comparison.

We consider the performance of these methods in a simulation study, where costs are assumed to follow Normal or Gamma distributions, to have positive and negative correlation with health outcomes, the instrument is strong (30% non-compliance) or weak (70% non-compliance), and the sample size, moderate (n=1,000) or small (n=100). We find that the proposed IV methods generally perform well. For example, in scenarios with Normally distributed cost data and a strong instrument, each method reports unbiased estimates. However, in these settings the 2SLS approach reports levels of Confidence Interval (CI) coverage that are above (positive correlation) and below (negative correlation) nominal levels. By contrast both the 3SLS and Bayesian methods report CI coverage close to nominal levels.

A structural Cox model for the causal effect of exposure among the exposed in an IV setting

Ditte Nørbo Sørensen[1], Torben Martinussen[1], Stijn Vansteelandt[2], Eric Tchetgen Tchetgen[3].
[1]*University of Copenhagen,*
[2]*Ghent University,*
[3]*Harvard University.*

In the simple instrumental variable (IV) setting with continuous variables and linear relations the causal effect can be found by two stage least squares. However when the relationship is nonlinear as is the case of survival outcome, there is a need for further development to do the proper IV analysis. The use of counterfactuals and structural models provides a setup where you can model the relationship of interest without specifying a joint model for the instrument, exposure, outcome and the unobserved confounder. This allows a relationship between the counterfactual survival function and the observed survival function that is the well-known Cox model. The focus will be on the case with one causal parameter, and some estimating equations are presented. The

method is directly applicable for binary exposure and instrument. In the more general case it involves specifying a model for the observable data, which can lead to situations where there is no solution to the estimating equations.

Relaxing the assumptions of causal discovery algorithms

Ricardo Silva and Robin Evans
University College London and University of Oxford

We consider the problem of estimating the average causal effect of a binary treatment X on an outcome Y. The faithfulness assumption says that conditional independence constraints in observational data warrants conclusions about causal structure, including a test of whether a given set of covariates is enough to block all unmeasured confounding between X and Y. This assumption can be justified by observing that its failure is akin to some sort of "effect cancellation" from multiple pathways in a causal graphical model, an event that seems unlikely. However, in practice, any statistical procedure aimed at detecting independence constraints with finite samples may be tricked by "near-cancellations" that are not ruled out by the faithfulness assumption, but which lead to accepted independence constraints. To account for that, we have recently introduced a procedure, the Witness Protection Program, which modifies existing causal discovery algorithms by allowing some degree of "path cancellation". This can be interpreted as an relaxation of faithfulness up to some hyperparameters that need to be chosen. The choice of such hyperparameters may not be straightforward in some applications. We propose a procedure for hyperparameter selection based on two principles: 1. assumptions that link observed dependencies to dependencies due to hidden confounders; 2. assessing the degree of failure of faithfulness by how coherent a causal discovery algorithm is at identifying the causal effect using different sets of constraints.

Causal and statistical inference with social network data: Massive challenges and meager progress

Elizabeth L. Ogburn
Johns Hopkins University

Interest in and availability of social network data has led to increasing attempts to make causal and statistical inferences using data collected from subjects linked by social network ties. But inference about all kinds of estimands, from simple sample means to complicated causal peer effects, is challenging when only a single network of non-independent observations is available. There is a dearth of principled methods for dealing with the dependence that such observations can manifest. We demonstrate the dangerously anticonservative inference that can result from a failure to account for network dependence and describe a few different avenues towards valid statistical and causal inference using social network data.

From Statistical Evidence to Evidence of Causality

Philip Dawid[1], Monica Musio[2], Stephen Fienberg[3]
[1]University of Cambridge,

[2]University of Cagliari,
[3]Carnegie Mellon University.

Statisticians typically study the `effects of causes' (EoC), but Lawyers and the Courts are more concerned with understanding the `causes of effects' (CoE). While EoC can be addressed using experimental design and statistical analysis, it is less clear how to incorporate statistical or epidemiological evidence into CoE reasoning, as might be required for a case at Law. Some form of counterfactual reasoning appears unavoidable, but this typically yields `answers' that are sensitive to arbitrary and untestable assumptions. It is however often possible to use statistical data to set bounds within which any answer must lie. With less than perfect data, these bounds will themselves be uncertain, leading to a compounding of different kinds of uncertainty. We illustrate these points with a Bayesian analysis of a case study in child protection.

Treatment effect of the treated: Understanding time-dependent confounding in survival analysis.

Odd O. Aalen[1], Jon Michael Gran[1], Rune Hoff[1], Kjetil Røysland[1]
[1]Oslo Centre for Biostatistics and Epidemiology, University of Oslo

The marginal structural model has been established as an important tool for analyzing survival data with time-dependent confounding. We propose an alternative based on analyzing the "treatment effect of the treated". This is a two-step procedure where we first use a missing data approach to estimate the values of time dependent covariates that a treated population would have were it not treated. For this we apply Farewell's linear increments model. We then combine this with an additive hazards model to estimate a causal estimate for the treatment effect of the treated. The method is illustrated by simulations and applications to real data. The results are compared to those from a marginal structural model.

We shall also illustrate and compare the marginal analyses with treatment-effect-of-the-treated analyses in a Markov model where explicit expressions can be given. The Markov model has the advantage of showing explicitly the presence of time-dependent confounding.

Marginal comparisons may give quite different numerical estimates from treatment-effect-of-the-treated analyses. In practice the latter type of analysis makes considerable sense, because the evaluation of a medication depends on the actual selection of patients for treatment. Also, in such a setting one does not need a complete positivity assumption.

A simple to implement algorithm for natural direct and indirect effects in survival studies with a repeatedly measured mediator.

Susanne Strohmaier[1], Nicolai Rosenkranz[2], Jørn Wetterslev[3] and Theis Lange[4]
[1]Oslo Centre for Biostatistics and Epidemiology, University of Oslo
[2]Department of Intensive Care, Copenhagen University Hospital
[3]Copenhagen Trial Unit, Centre for Clinical Intervention Research, Copenhagen University Hospital
[4]Department of Biostatistics, University of Copenhagen

Important questions within the fields of social sciences, epidemiology as well as clinical trial research involve the challenge of decomposing the effect of an interventions into direct and indirect effects working through a defined mediator, thereby aiming for a better understanding of the underlying mechanisms. For the case of a single and multiple mediators measured at a single point in time, researchers have established theoretical properties and developed practical tools for the analysis of a broad class of mediator and outcome models (e.g. Lange et al. (2012, 2014)) by employing the counterfactual framework. However, data structures are often more complex than the described scenarios.
We present an extension to the procedure by Lange et al. to the setting of a time-to-event outcome and a repeatedly measured mediator, where the number of measurements is determined by survival time. We suggest an estimation algorithm, that allows for direct parametrisation of direct and indirect natural effects and is easy to implement using standard software. The proposed method enables us to analyse the mediating role of KDIGO (a measure of severity of kidney impairment) on mortality in the Scandinavian Starch for Severe Sepsis/Septic Shock trial (6S) comparing two substances for fluid resuscitation among patients with severe sepsis admitted to intensive care units.

Mediation analysis of randomised experiments.

Sjouke Vandenberghe and Stijn Vansteelandt
Department of Applied Mathematics, Computer Sciences and Statistics, Ghent University

In this talk, we will make use of mediation analysis to decompose the effect of a randomised treatment on an end-of-study outcome into its indirect effect via a certain mediator and the remaining direct effect. This is motivated by a re-analysis of the EORTC 10994/BIG 1-00 randomised phase 3 trial, which aims to infer the indirect effect of taxane and anthracycline based chemotherapies on overall survival of advanced breast cancer patients, that is mediated by pathological complete response. We improve the efficiency of estimators of natural direct and indirect effects that were previously proposed by Tchetgen Tchetgen (2011) and Tchetgen Tchetgen and Shpitser (2012), by exploiting the fact that the data originate from a randomised experiment. The resulting estimators are less model-dependent than popular estimators based on the mediation formula: they do not demand correct specification of the model for the mediator, because they rely on the known randomisation probabilities. Results from a simulation analysis are shown, in which the proposed estimators are compared with competing estimators in terms of efficiency and bias in a variety of settings, for example under misspecification of the mediator model. Data analysis results will be discussed for the EORTC 10994/BIG 1-00 randomised phase 3 trial.

Bounds on biological/causal interactions

Arvid Sjölander
University Stockholm

A common goal of epidemiologic research is to study the interaction between two exposures. A potential source of confusion is that the term “interaction” has quite different interpretations in different fields. In statistics, the term is typically used as a synonym for “product terms” between the two exposures in a statistical (e.g. linear or logistic) model. In applied fields, such as genetic epidemiology, the term is often used to describe a physiological or biological mechanism in which the two exposures (e.g. genetic alleles) “cooperate”. The former is often labelled “statistical” interaction, and the

latter has been referred to as “biological” or “causal” interaction. Rothman and Greenland (Modern Epidemiology, 2nd ed) proposed a formal definition of biological interaction, based on potential outcomes. For binary exposures, they showed that additive statistical interaction implies biological interaction. This means that the presence of biological interaction can be tested, i.e. we can say whether biological interaction is present or not. However, they also showed that the magnitude of biological interaction can not be estimated, unless one is prepared to make very strong and generally unrealistic assumptions. In this seminar we demonstrate that biological interaction can be bounded, i.e. it is possible to provide an upper and lower limit for its magnitude. We present such bounds in a very general setting, allowing for categorical exposures with arbitrary many levels. We demonstrate that the bounds can be made significantly tighter under an assumption of monotone exposure effects. We present an application of the bounds to a study of gene-gene interaction in rheumatoid arthritis.

Causal Interaction in High Dimension

Kosuke Imai
Princeton University

Estimating causal interaction effects is essential for the exploration of heterogeneous treatment effects. In the presence of multiple treatment variables with each having several levels, researchers are often interested in identifying the combinations of treatments that induce large additional causal effects beyond the sum of separate effects attributable to each treatment. We show, however, the standard approach to causal interaction suffers from the lack of invariance to the choice of baseline condition and the difficulty of interpretation beyond two-way interaction. We propose an alternative definition of causal interaction effect, called the marginal treatment interaction effect, whose relative magnitude does not depend on the choice of baseline condition while maintaining an intuitive interpretation even for higher-order interaction. The proposed approach enables researchers to effectively summarize the structure of causal interaction in high-dimension by decomposing the total effect of any treatment combination into the marginal effects and the interaction effects. We also establish the identification condition and develop an estimation strategy for the proposed marginal treatment interaction effects. Our motivating example is conjoint analysis where the existing literature largely assumes the absence of causal interaction. Given a large number of interaction effects, we apply a variable selection method to identify significant causal interaction. Our analysis of a survey experiment on immigration preferences reveals substantive insights the standard conjoint analysis fails to discover. Joint work with Egami Naoki.

The paper is available at http://imai.princeton.edu/research/int.html

Bias-Reduced Doubly Robust Estimation.

Karel Vermeulen and Stijn Vansteelandt
Ghent University

Doubly robust (DR) estimators are asymptotically unbiased when either an outcome regression model or a propensity score model is correctly specified, regardless of which. While theoretically appealing, DR estimators have been the subject of recent debate. The reason is that model misspecification is likely to affect all working models in practice, and thus the very premise that at least one of both working models is correctly

specified, lives on shaky grounds. Moreover, the performance of DR estimators can be sensitive to the choice of estimators used for fitting the working models, and can sometimes be worse than that of competing estimators that do not enjoy the double protection property.
In this talk, I will show that, interestingly, some DR estimators partially retain their robustness properties even under misspecification of both working models. In particular, I will propose a simple and generic estimation principle for the nuisance parameters indexing each of the working models, which we call bias-reduced DR estimation (Vermeulen and Vansteelandt, 2015). It is designed to improve the performance of the DR estimator of interest, relative to the default use of maximum likelihood estimators for the nuisance parameters by locally minimizing the squared first-order asymptotic bias of the DR estimator under misspecification of both working models. We discuss the basic proposal, which is based on parametric models for the nuisance models, as well as extensions that employ machine learning algorithms. Simulation studies confirm the desirable finite-sample performance of the proposed estimators relative to other proposals.

Doubly-robust dynamic treatment regimen estimation via weighted least squares.

Michael P Wallace, Erica E M Moodie and David A Stephens
McGill University

Personalized medicine is a rapidly expanding area of health research wherein subject level information is used to inform treatment. Dynamic treatment regimens (DTRs) are one means by which personalized medicine can be studied theoretically and applied in practice. DTRs are sequences of decision rules which take subject information as input and provide treatment recommendations as output. Such regimens therefore tailor each treatment decision to a patient's unique circumstances, but can also identify management plans which optimize long-term outcomes by accommodating potentially obscure delayed treatment effects and other complex interactions. However, taking such factors into account can complicate the problem of causal inference in this context. We consider the blip: a structural nested mean model of the expected difference in the (potentially counterfactual) outcome when using a baseline treatment instead of observed treatment. DTR estimation in this context therefore relies on estimating blip parameters and numerous methods have been proposed for this purpose. We present a new approach which uses standard weighted ordinary least squares regression to control for the potentially confounding effects of covariate-dependent treatment. This builds on two established methods: Q-learning and G-estimation, offering the doubly-robust property of the latter but with ease of implementation akin to the former. We outline the underlying theory and demonstrate the double-robustness and efficiency properties of our approach through illustrative examples. Finally we discuss model assessment, demonstrating diagnostic plots for our method, and how the double robustness property itself may be leveraged to investigate model validity.

Tackling non-ignorable dropout in the presence of time-varying confounding.

Marco Doretti[a], Sara Geneletti[b], Elena Stanghellini[a]
University of Perugia[a],
London School of Economics and Political Science[b]

We propose a slight extension of the g-computation algorithm (Robins 1986, Ryan et al. 2012) to explore the sensitivity of time-varying confounding adjusted estimates to different dropout mechanisms. We implement the Heckman correction for a general number of occasions and explore selection models to investigate situations where the dropout process is driven by unobserved variables and the outcome respectively. The analysis is embedded in the Bayesian framework, which provides a number of advantages. These include fitting a hierarchical structure to processes that repeat over time and avoiding exclusion restrictions in the case of the Heckman correction. Our approach to causal inference relies on the Decision Theoretic framework (Dawid and Didelez, 2010), which highlights the need of an additional assumption we term *No regime dropout dependence* (NRD). Our methods are applied to data from the Counterweight Programme pilot, a UK protocol to address obesity in primary care (Taubman et al. 2009). A simulation study evaluating the performance of the proposed estimators is also presented.

Estimating the optimal treatment sequence for graft-versus-host-disease following bone marrow transplantation

Erica E.M. Moodie
McGill University

Determining the optimal sequence of treatments is an area of active research in the statistical, computer sciences, and medical communities. Much of the methodological work done to date has focused on the continuous outcome settings. While extensions to time-to-event settings have been considered, the methods do not conform to a wide variety of problems in which we would like to optimize treatment sequences. Take for example, the role of immunosuppressive therapy that depletes T cells in prophylaxis and treatment of graft-versus-host-disease (GVHD). To address the question of whether the sequence of administration of non-specific, highly T-lymphodepleting therapies in GVHD prophylaxis and in treatment of refractory GVHD impacts survival, and to identify donor-related and patient-related factors that may guide individualized selection of the classes and sequence of immunosuppressive agents over time, we develop a parametric Q-learning approaches, and apply them to a large cohort drawn from a bone marrow transplantation registry from 1995 to 2007.

Causal inference in survival analysis: An example from prostate cancer.

Kjetil Røysland
University of Oslo

What makes survival analysis slightly different from much other statistics is that we almost always have to deal with censoring. Still, we are interested in parameters we would see if the censoring had been completely prevented. This, however, is really a claim about causation, and can be treated formally using graphical models and techniques from causal inference.
We present an analogy to the back-door criterion, based on local independence graphs and delta-separation, that allows us to determine if given parameters would take the same value as if the censoring was prevented, even if the actual data were censored. We suggest a method for achieving estimates of such parameters based on continuous-time censoring weights and stochastic differential equations, driven by cumulative hazards. Unlike the usual discrete-time censoring weights, their continuous-time counterparts enable consistent estimates of several common parameters in survival analysis. This can

be seen by applying standard techniques from the theory of stochastic differential equations. To illustrate the methodology, we provide an example that compares the effects of radiation therapy and radical prostatectomy on Norwegian patients with prostate cancer.

Acute health shocks and labour market exits.

Andrew M. Jones[1], Nigel Rice[1,2], Francesca Zantomio[3]
[1]Department of Economics and Related Studies, University of York
[2]Centre for Health Economics, University of York
[3]Ca' Foscari University of Venice

The financial consequences of early labour market exit can be substantial and long-lasting. This paper investigates the labour supply response to acute health shocks defined by the incidence of cancer, stroke, or heart attack, for working age individuals in the UK; a group rarely considered in previous studies which generally focus on older individuals. We draw on data from Understanding Society which offers a unique combination of a large sample, a panel dimension together with a broad range of socio-demographic, health and labour market information that makes it particularity suited to this study. Our identification strategy exploits uncertainty in both the occurrence and timing of an acute health shock. We follow individuals until they experience either a first occurrence of a health shock, or a re-occurrence. We compare labour supply responses to those observed in a control group. Controls are defined through a combination of coarsened exact matching and parametric propensity score estimation. The panel dimension of the data allows us to condition on unobserved individual heterogeneity. Our results indicate that, on average, experiencing an acute health shock significantly reduces labour market participation, with a stronger response to an additional, as opposed to a first, shock. In general younger workers of both genders display a stronger labour market attachment than older counterparts conditional on a health shock. Older and more educated women exhibit the strongest retraction despite experiencing less disabling shocks. This suggests an important role for preferences, financial constraints, and intra-household division of labour in explaining labour supply adjustments.

Causal inference in a multi-state model for sickness absence and return to work.

Jon Michael Gran[1], Stein Atle Lie[2], Irene Øyeflaten[3,4], Odd O. Aalen[1]
1. Oslo Centre for Biostatistics and Epidemiology,
Department of Biostatistics, University of Oslo, Norway.
2. Department of Clinical Dentistry, University of Bergen, Norway.
3. National Centre for Occupational Rehabilitation, Norway.
4. Uni Health, Uni Research, Norway.

Multi-state models, as an extension of traditional models in event history analysis, provide a flexible framework for analysing multiple transitions in a unified manner. In this paper we study three different approaches for causal inference in multi-state models. One approach is to estimate the effects of interventions in terms of manipulating transition intensities between states, while the two other approaches are based on inverse probability weighting and g-computation.
The methods are applied to registry data on social security benefits, which is a valuable source for analysing sick leave, disability and employment. Recent work on data from

Norwegian and Danish registries has proven multi-state models to be very suitable framework for analysing this kind of data. Individuals incapable of working due to disease or injury typically move on and off different benefits with time, which naturally can be perceived as moving between a given set of states in a multi-state model.
In this paper we analyse national registry data on sick leave benefits, work assessment allowance or disability pension from the Norwegian Labour and Welfare Administration coupled with data from a multicenter cohort of individuals participating in work-related rehabilitation programs. The detailed covariate information is used to predict future return to work and disability for such participants, and to compare the effect of treatment regimes, such as the use of full versus part time sick leave.

Stable Weights that Balance Covariates for Causal Inference and Estimation with Incomplete Outcome Data

José R. Zubizarreta,
Columbia University

Weighting methods that adjust for observed covariates, such as inverse probability weighting, are widely used for causal inference and estimation with incomplete outcome data. Part of the appeal of such methods is that one set of weights can be used to estimate a range of treatment effects based on different outcomes, or a variety of population means for several variables. However, this appeal can be diminished in practice by the instability of the estimated weights and by the difficulty of adequately adjusting for observed covariates in some settings. To address these limitations, this paper presents a new weighting method that finds the weights of minimum variance that adjust or balance the empirical distribution of the observed covariates up to levels prespecified by the researcher. This method allows the researcher to balance very precisely the means of the observed covariates and other features of their marginal and joint distributions, such as variances and correlations and also, for example, the quantiles of interactions of pairs and triples of observed covariates, thus balancing entire two- and three-way marginals. Since the weighting method is based on a well-defined convex optimization problem, duality theory provides insight into the behaviour of the variance of the optimal weights in relation to the level of covariate balance adjustment, answering the question, how much does tightening a balance constraint increases the variance of the weights? Also, the weighting method runs in polynomial time so relatively large data sets can be handled quickly. An implementation of the method is provided in the new package sbw for R. This paper shows some theoretical properties of the resulting weights and illustrates their use by analysing both a data set from the 2010 Chilean earthquake and a simulated example.

A formal treatment of Regression Discontinuity Designs.

Nayia Constantinou[1] and Aidan G. O'Keeffe[2]
School of Mathematics, University of Bristol, UK
Department of Statistical Science, University College London, UK

Regression Discontinuity Designs (RDDs) occur when a decision to apply treatment is linked to some continuous `assignment variable' through a decision rule. Typically, such decision rule will specify treatment if the subject's associated assignment variable value lies at or above a pre-specified threshold and will specify no treatment if the subject's corresponding assignment variable lies below the threshold. Under the assumption that

subjects with similar assignment variable values are exchangeable and that the decision rule is identical for all subjects, a comparison of the response variable values between individuals who lie at or above the threshold and individuals who lie below the threshold may be considered appropriate for the calculation of a causal effect of the treatment on the outcome of interest. We take a rigorous Decision Theoretic approach to study two versions of such designs, the strict RDD and the fuzzy RDD, and formally explore conditions which allow us to make causal inference. We illustrate our analysis using a real dataset that involves the prescription of statins based on the cardiovascular risk score.

Regression discontinuity designs: The challenge of binary outcomes.

Sara Geneletti[1], Aidan O'Keeffe[2], Gianluca Baio[2]
1. London School of Economics and Political Science,
2. University College London

The regression discontinuity design (RDD) is a natural experiment that exploits the fact that many treatments are assigned according to pre-defined rules. An example is the prescription of statins, a class of cholesterol-lowering drugs. Individuals are prescribed statins if their risk of developing CVD in the subsequent 10 years, as calculated by an appropriate risk calculator, exceeds 20. If we can plausibly assume that individuals within a certain distance of the threshold belong to the same population with respect to the characteristics that inform the assignment rule and determine the outcome, then the threshold can be seen as a quasi-random intervention which assigns the treatment to those that are just above the threshold and assigns no treatment to those that fall just below the threshold. We can then exploit this random assignment to estimate the (causal) effect of the treatment for individuals in the region around the threshold. This can be straight-forward in the case of continuous outcomes; however, binary outcomes represent a new set of challenges. These are methodological: causal estimates based on instrumental variable methods for binary outcomes require very strong and often implausible assumptions, and also practical: defining a meaningful binary outcome and choosing prior distributions amongst others. In this presentation we talk about these issues and apply out methods to a dataset on statin prescription in primary care.

Dynamic Causal Inference for a binary outcome in a Regression discontinuity design using local independence.

Aidan G. O'Keeffe
Department of Statistical Science, University College London, UK

Regression discontinuity (RD) designs have been developed as method for causal inference from observational data, through the exploitation of a naturally occurring intervention rule. Most RD design methods have focussed on causal inference where the outcome of interest is continuous. Scenarios in which the outcome of interest is binary have received less attention.
A regression discontinuity design consists of three processes: an intervention process, a binary outcome process and a threshold attainment process (where the decision on whether or not to apply an intervention is taken based on the attainment of a fixed 'threshold' by a subject-specific continuous 'assignment variable'). These three processes may change at different points over time and we examine the dynamic relationships amongst the processes and consider how associated causal relationships may be inferred.

We use properties of local independence to develop a suitable framework in which the dependencies amongst these processes may be modelled over time, under particular assumptions, and causal inference made concerning the effect of a previous intervention on the time until the occurrence of a binary outcome. We demonstrate the developed methods using simulated data and present a short example using real data on the prescription of statins for the primary prevention of cardiovascular disease in UK primary care.

Program Evaluation with High-Dimensional Data

Victor Chernozhukov
Massachusetts Institute of Technology

We consider estimation of policy relevant treatment effects, such as average treatment effects or quantile treatment effects in an environment where there may be many more control variables available than there are observations. We are specifically interested in settings with heterogeneous treatment effects and endogenous receipt of treatment. We impose that key reduced form predictive relationships are approximately sparse to make informative inference possible which allows estimation and inference to proceed by selecting an appropriate set of control variables. To accommodate estimation of a wide variety of treatment effects estimators, we provide a number of new general results that are used to establish good estimation properties and uniformly valid post-selection inference for a continuum of target parameters defined by moment restrictions that may depend on a continuum of high-dimensional nuisance functions. These results provide a general set of conditions under which inferential results for function-valued parameters of interest following model selection are uniformly valid over a wide range of models which are applicable outside of the treatment effects context and of independent interest. We illustrate the use of the proposed treatment estimation methods with an application to estimating the effect of 401(k) participation on accumulated assets. Relevant papers: http://arxiv.org/abs/1311.2645, http://arxiv.org/abs/1201.0224

Instrumental variables and survival analysis

Els Goetghebeur (joint with Jozefien Buyze)
Ghent University

Estimation of the causal exposure effect on a survival outcome from observational data must account for possible confounding. When important confounders remain unmeasured, instrumental variables can offer a solution. For an uncensored continuous outcome, two-stage least squares and semi-parametric structural mean models are often used. We extend the first method assuming a log-linear model for the right-censored survival outcome while accounting for non-informative censoring and compare it with the structural accelerated failure time approach in our setting. We are guided by 2 case studies: quality of care assessment over hospitals and the analysis of non-compliance in an HIV prevention trial.

We find that for the categorical exposure, we need stronger conditions on the instrumental variables when fitting separate models comparing each exposure level to all other levels combined then when fitting a single model comparing each exposure level to a common 'reference level'. We saw negligible bias on exposure effects and

improved efficiency when we do not re-censor for the AFT approach. Once the causal effect of exposure has been estimated, we proceed to estimate and compare derived survival curves under different potential exposures from back-transformed observed outcomes.

We end with a discussion of the bias efficiency trade off under the different current options, referring also to the structural proportional hazard s model and considering the impact of correlation between multiple instruments.

This presents joint work with Jozefien Buyze.

POSTER ABSTRACTS

On the advantages of threshold blocking.

Fredrik Sävje
Department of Economics, Uppsala University.

A common method to reduce the uncertainty of causal inferences from experiments is to assign treatments in fixed proportions within groups of similar units—blocking. Previous results indicate that one can expect substantial reductions in variance if these groups are formed so to contain exactly as many units as treatment conditions. This approach can be contrasted to threshold blocking which, instead of specifying a fixed size, requires that the groups contain a minimum number of units. In this paper I will investigate the advantages of respective method. In particular, I show that threshold blocking is superior to fixed-sized blocking in the sense that it, for any given objective and sample, always finds a weakly better grouping. For blocking problems where the objective function is unknown, this need, however, not hold and a fixed-sized design can perform better. I specifically examine the factors that govern how the methods perform in the common situation where the objective is unconditional variance but groups are constructed based on covariates. This reveals that the relative performance of threshold blocking increases when the covariates become more predictive of the outcome.

Application of Marginal Structural Models to unbalanced longitudinal health data.

Edmore Chamapiwa[1], David Reeves[1], Darren Ashcroft[2], Evangelos Kontopantelis[1] and David Springate[1]
[1]Centre for Primary Care, University of Manchester, Manchester
[2]Manchester Pharmacy School, University of Manchester

Background: Marginal Structural Models (MSMs), a class of structural causal models, are being increasingly used in the analysis of complex longitudinal health data because of their ability to give unbiased estimates of a time-varying treatment in the presence of time-varying confounding/mediating covariates. However, MSMs assume that observations occur at regularly separated time points for all patients, whereas in "real-life" health record data, different patients are commonly seen and measured at different and irregular time points. The impact of unbalanced, but more realistic, data on the performance of MSMs is unknown.
Objective: To evaluate performance of inverse-probability-weighted MSMs in unbalanced longitudinal data
Methods: A simulation study was conducted to compare treatment effect estimates from inverse-probability-weighted MSM, unadjusted generalised estimating equation (GEE) model and adjusted GEE model. Unbalanced longitudinal data was generated by sampling

Principal stratification, treatment evaluation and robustness of sample selection models-simulation study, with application to job corps data.

Giuseppe Moscelli[1,2,] – Roberto Rocci[1,3]*
1: "Tor Vergata" University of Rome
2: Centre for Health Economics (CHE), University of York
3: LUISS University, Rome

We compare the statistical performance of several estimators of the Principal Stratification and Heckman selection models in a large Monte Carlo study. The aim is the unbiased and efficient estimation of the effect of an exogenous and monotone treatment. Compared to the previous literature, we test the robustness of the different estimators to misspecifications due not only to the data generating process, but also the regression functional form and the distribution of the error terms. Moreover, we formally consider both cases, with and without an exclusion restriction. We propose a parametric modification of the classic Principal Stratification model exploiting the monotonicity assumption, and a semi-parametric modification of the classic Principal Stratification model that is more robust to regression functional form misspecifications. We show that Principal Stratification estimators are generally less sensitive than Heckman estimators to different sources of misspecification, especially when an exclusion restriction is not available. The semi-parametric Principal Stratification extension proves to be frequently less biased and more efficient than the classic Principal Stratification and Heckman models for the estimation of the treatment effect of interest. We provide an empirical application using data from Job Corps, one of the largest training programs in the U.S. We find a positive effect of the participation to the training program on income. The use of the semi-parametric Principal Stratification extension adds efficiency gains compared to the classic estimator when we include a covariate potentially affecting nonlinearly both the endogenous selection into employment and income.

Structure makes function makes structure.

Yoli Shavit and Pietro Lio'
Computer Laboratory, University of Cambridge, Cambridge, CB3 0FD, UK

An important paradigm in proteomics is 'structure makes function'. High-throughput Chromosome Conformation Capture (3C) technologies, detecting contact frequencies between genomic segments, have made it possible to test this paradigm for genome architecture. However, despite the large availability of Omics data, such as methylation Histone modification and gene regulation, and of High-throughput 3C data, such as 5C and Hi-C, there are currently limited methodologies for putting together spatial and functional genomic information. Such integration is especially challenging due to differences in resolution between data sets. We present a unified framework for the analysis of 3C and multi-omic data, taking advantage of a multi-scale (fractal) model of genome packing. We show that our approach could provide insights into the complex structure of chromosomes and present derived metrics for testing for causality between spatial and omics data. This approach could further prove useful for studying how structure and function at the nucleus act and react together with respect to external signals, and how this in turn translates into phenotype and disease conditions.

Mediation analysis of time-to-event outcomes based on pseudo-observations.

Theis Lange
Section of Biostatistics, University of Copenhagen

When considering time-to-event outcomes the most used effect measures are based on hazards, in particular hazard ratios. However, it is by no means clear that hazards ratios are always the most appropriate effect measure. For instance will restricted mean life times be easier to communicate and interpret in oncology trials with a high mortality. When doing mediation analysis this is even more acute as measures such as the mediated proportion based on hazard ratios are very difficult to communicate (correctly) to the non-causal inference community. In this talk I propose to use pseudo-observations computed from censored time-to-event data in combination with natural effects models to obtain measures of mediation based on restricted mean life times. The method is implemented by combining existing R packages for computing pseudo observations with the just released medflex package. The medflex package allows natural effects models to be estimated without doing "customized" coding. Thus the proposed pseudo-observation based method can be implemented using existing functions. The talk will also briefly mention the capabilities of the medflex package.

Applying optimal mediation methods to clinical RCT data: Example from a low back pain intervention study.

Rebecca Case, Gemma Mansell, Daniëlle van der Windt
Research Institute for Primary Care and Health Sciences, Keele University

Randomised controlled trials (RCTs) of psychosocial interventions to reduce disability in low back pain populations have often been found to have only small mean effects. Identifying explanations for small effect sizes can be an important step in improving future trials. We provide an example of how mediation analysis can be used to help explain small effect sizes in a cluster RCT which compared a minimal psychosocial intervention with usual care by the general practitioner for patients with acute or subacute low back pain. Our mediation analysis used the causal inference approach (through the -paramed- command in Stata) to investigate several variables measured in the trial (catastrophising, fear-avoidance beliefs, distress, rest, increased exercise and staying active) that were hypothesised to have an indirect effect on the effect of treatment on outcome (physical function). Adjustments were made for potential confounding variables (age, sex and baseline measures of the potential mediators and outcome). While a significant direct effect was found between treatment and several potential mediator variables (fear-avoidance, increased exercise and staying active), only fear-avoidance beliefs had a significant indirect effect (-0.38, 95% CI -0.86 to -0.05) and this was no longer significant in the adjusted analysis (-0.20, 95% CI -0.59 to 0.04). The aim of this study was to not only investigate the reasons for small effects in back pain trials, but also to discuss the potential benefits and pitfalls of using this particular mediation approach compared to other approaches and make some preliminary recommendations for applying mediation analysis to clinical studies.

Obesity paradox: Manifestations and Explanations.

J Candlish, E Badrick, A Reehan, M Sperrin.
University of Manchester

An obesity paradox has been reported among patients with certain diseases which suggests that overweight and obese individuals have lower mortality than normal weight individuals. Several explanations have been hypothesised for this paradoxical association including reverse causality, different cohorts of people developing the disease, and collider stratification bias. Conditioning on presence of disease (collider) may induce false correlations or strengthen correlations between risk factors (BMI) and outcome (mortality), because of unmeasured confounders. A consensus on the explanations for the obesity paradox has yet to be reached.
This work investigates the obesity paradox in the context of Type 2 diabetes mellitus (T2DM) through a review of the different causal structures that may be responsible for the paradox, an analysis of individuals from Salford, UK with T2DM to represent how the paradox manifests itself. Sensitivity analysis were undertaken using the Salford T2DM cohort altering inclusion and exclusion criteria, producing varying results suggesting possible reasoning behind the lack of consensus in the literature assessing the paradox among T2DM. An extension of the DAG framework representing collider bias will be presented, based on the hypotheses of two 'types' of diabetes. We will present simulations that assess the extent of the biases caused under different scenarios of confounding and selection bias. Finally, we will discuss the wider issue of conditioning on a disease state affected by exposure and sharing common causes with the outcome, of which the obesity paradox is just one example.

Mediation and moderated mediation: Accounting for confounding and counterfactual thinking: A comparison of 8(+2) approaches.

George Chryssochoidis
Norwich Business School, University of East Anglia

Recall mediation and moderating modeling? Mediation (henceforth '*Med*') refers for instance to what is depicted as 'Model 4' in Hayes (2013) and moderated mediation (henceforth '*ModMed*') refers to what is depicted as 'Model 59' in Hayes (2013) (or 'total effect' moderation model in Edwards and Lambert, 2007:4). Business research is at cross-roads. Lack of considering confounding or lack of causal thinking are now seen as having substantially biased past research. 'Confounding' has been defined primarily as non-modelling model-relevant variables ('confounders') (VanderWeele and Shpitser, 2013) resulting to inaccurate estimates (Antonakis *et al.*, 2010). Moreover, causal reasoning (unlike associational reasoning mostly practiced under a SEM framework) involve a notion of how the world would have been had should an element in the Med or ModMed model been different. In doing so, the definition of direct and indirect effects involve quantities that are not all observable: Y(x): the potential values of Y that would have occurred had X been set, possibly counter to fact, to the value x; M(x): the potential values of M that would have occurred had X been set, possibly counter to fact, to the value x. Similarly for Y(x, m) and Y(x, M(x*)).
Research–especially in economics and epidemiology- has proposed however methods to deal with confounding and also causal (counterfactuals) effects estimation. These include for instance: a) Instrumental variables-based (e.g., Garen, 1984); b) Strata based (Emsley *et al.,* 2010); c) Propensity score matching (Guo and Frazer, 2010); d) VanderWeele's 4-way effect decomposition (VanderWeele, 2014); e) Paramed (Emsley and Liu, 2013); f) Imai *et al.* (2010a; b; 2011) approach; g) Muthen and Asparouhov (2014) approach; h) Latent IV (Ebbes et al., 2009) approach.
However, no comparative study exists identifying the (dis)similarities of the outcomes produced by these alternative suggested solutions. The objective of this presentation is to compare these approaches using the same dataset (drawn from management) and

contrast the extent and type of direct and indirect effect they produce. The focus is on cross-sectional data.

Bounds for the Probability of Causation in the Presence of a Mediator.

R. Murtas[1], M. Musio[1] and A. P. Dawid[2]
1. University of Cagliari
2. University of Cambridge

An individual has been subjected to some exposure and has developed some outcome. Using data on similar individuals, we wish to evaluate, for this case, the probability that the outcome was in fact caused by the exposure. Even with the best possible experimental data on exposure and outcome, we typically can not identify this "probabiity of causation" exactly, but we can provide information in the form of bounds for it. Under appropriate assumptions, these bounds can be tightened if we can make other observations (e.g., on non-experimental cases), or measure additional variables (e.g., covariates). In this work we develop such improved bounds for the case that a third variable mediates the effect of exposure on outcome.

Addressing unmeasured confounding in health interventions using observational data: A systematic review.

Adam J. Streeter, Nan X. Lin, Louise Crathorne, Alessandro Blé, Morwenna Rogers, Marcel Haasova, Chris Hyde, David Melzer, William Henley
University of Exeter

Introduction: There is growing interest in medical 'big data' – millions of patient records from administrative databases that dwarf data from clinical trials, and provide opportunities to study interventions in real-world populations, among a wider variety of risk groups. However, such data sources were not designed with causal inference in mind, and many results are likely to be confounded and biased by differences in treatment group characteristics. Furthermore complete information on these confounders may not be available. The influence of unmeasured confounding can be assessed through sensitivity analysis, and adjustment achieved using some analytical methods, depending on the tenability of their assumptions.
Aims: Our systematic review examined applications of sensitivity analyses and methods to adjust for unmeasured confounding in longitudinal, observational health data.
Results: Of the 520 citations returned by our search, 52 were eligible for review. The predominant approach in 32 papers was through instrumental variable analysis. Three used propensity score calibration. Explicitly longitudinal methods, including difference-in-differences, regression-discontinuity, the prior-event rate ratio method and a hybrid time-series, were evident in fewer than four papers each. Unmeasured confounding was characterised through sensitivity analysis in 15 of the included papers.
Conclusions: The review suggested that either the methods did not fully exploit longitudinal information or have yet to be fully developed and disseminated. This may reflect that the longitudinal approach is not explicitly integrated into the current framework of causal inference. However promising methods do exist, and these could be further developed to provide tools for the current health-data 'goldrush'.

Matching combined with regression versus the Synthetic Control approach for evaluating treatment effects: A simulation analysis and case study.

Stephen O'Neill[1], Noémi Kreif[1], Richard Grieve[1] and Jasjeet S. Sekhon[2]
[1] Department of Health Services Research & Policy, LSHTM, University of London
[2] Department of Political Science and Department of Statistics, UC Berkeley, Berkeley

Difference-in-difference (DiD) estimators can provide unbiased estimates of treatment effects when time-invariant unobserved confounders have time-constant effects, i.e. the parallel trends (PT) assumption holds. The Synthetic Control (SC) approach allows for these aspects of unobserved confounding by weighting the control units to match the pre-treatment outcomes and characteristics of the treated unit(s). However, with multiple treated units, the SC weights may leave residual imbalances between the treatment groups. Genetic Matching (GM) uses an automated search algorithm to maximise balance across the whole distribution of observed covariates (e.g. pre-treatment outcomes). Combining Genetic Matching with DiD regression (GMDiD), can balance observed covariates, but also allow for aspects of unobserved confounding (as per SC).
Monte Carlo simulations were conducted to contrast the methods' performance in various settings including: where (1) covariates are balanced and PT holds; (2) observed covariates are balanced but PT fails and (3) observed covariates are imbalanced beyond the mean and PT fails. We report bias and root mean squared error (RMSE).
When PT holds, all methods report unbiased estimates of the treatment effect, with the DiD estimates having the lowest RMSE. When PT fails, SC provide estimates with low bias versus DiD, but high RMSE, reflecting the sparsity of the weights; relatively few control units received a positive weight. GMDiD reports estimates with low bias and uses outcome data from more control units, and so offers efficiency gains compared to SC. The methods were examined in a reanalysis of the impact of Best Practise Tariffs for Stroke in the UK.

Developing a chaining approach for estimating the lifetime returns of childhood health interventions.

Alex J Turner, Eleonora Fichera, & Matt Sutton
Manchester Centre for Health Economics, Institute of Population Health, University of Manchester

Background: The benefits of childhood health interventions should be assessed in terms of lifetime health consequences, but intervention studies often focus only on short-term changes. Previous studies have attempted to match trial data to birth cohort datasets in order to proxy outcome trajectories across the life-course. However, even the cohort datasets with the longest follow-up are incomplete and are inevitably dated.
Aim: To describe and demonstrate a method for estimating the full lifetime health returns to interventions in childhood via matching of data from intervention studies to short longitudinal datasets collecting data on all ages.
Method: Using Coarsened Exact Matching, child-level data on outcomes and other characteristics from an intervention study are matched to longitudinal data on children of the same age. Lifetime trajectories are generated by chaining one-year transitions from consecutive starting ages, using only the two most recent waves of longitudinal data. The trajectory for each individual ends when death occurs in a transition between waves. Confidence intervals are generated using bootstrapping.
Data: 989 children from a randomised controlled trial of a school-based social and emotional wellbeing intervention are linked to 790 children from a large longitudinal dataset, Understanding Society.

Results: The mean 0.8-unit improvement in the Strength and Difficulties Questionnaire score generated by the intervention is associated with an additional 0.9 lifetime Quality-Adjusted Life Years.
Implications: Coarsened Exact Matching between intervention studies and consecutive one-year transitions in longitudinal datasets offers a feasible method for estimating lifetime outcomes using the most up-to-date information on changes over the life-course.