How Drug Availability Affects Consumer Compliance: The Role of Market Concentration

Jiancai Liao[(1)], Chuhan Liu[(2)], Jian Ni[(2)], Haizhong Wang[(3)].

(1) Department of Marketing, Aalto University School of Business, P.O. Box 21220, FI-00076 Aalto, Finland,
(2) Department of Economics, Johns Hopkins University Carey Business School, 100 International Drive, Baltimore, USA,
(3) Department of Marketing, SunYat-sen University School of Management, No. 135, Xingang Xi Road, Guangzhou, China

Abstract: Medicine shortages have been a persistent issue affecting low, middle, and high-income countries. Many countries have developed various drug availability guarantee policies (hereafter, DAGP) to overcome the problem.
However, the behavioral consequences of introducing DAGP remain unclear. This article investigates the impact of drug availability on a crucial behavioral outcome: consumer compliance. We utilize the exogenous shock of the ”4+7” drug volume-based procurement policy, which directly influences drug availability, to establish a causal relationship between drug availability and consumer compliance using a difference-in-difference model.
Our findings indicate that increased drug availability negatively impacts consumer compliance levels, with the decline in compliance being particularly pronounced in low-concentration markets that offer numerous product alternatives. Drawing upon institutional theory, we argue that prevailing institutional logics within a specific institutional environment influence how actors respond to regulatory changes, such as implementing the ”4+7” drug volume-based procurement policy. Market concentration plays a pivotal role in shaping consumer sensitivity to the policy, as distinct institutional logics drive compliance in high-concentration markets while promoting exploration in low-concentration markets.

Keywords: Consumer compliance; Drug availability; Market concentration; Institutional theory.

SmaC: Spatial Matrix Completion Method

Giulio Grossi[(1)].

(1) Department of Statistics, Computer Science, Applications "G. Parenti", University of Florence, Italy

Abstract: Synthetic control methods are commonly used in panel data settings to evaluate the effect of an intervention. In many of these cases, the treated and control time series correspond to spatial areas such as regions or neighborhoods. We work in the setting where a treatment is applied at a given location and its effect can emanate across space. An area of certain size around the intervention point is assumed to be the treated area.
Synthetic control methods can be used to evaluate the effect that the treatment had in the specific area, but it is often unclear how far the treatment's effect propagates. Therefore, researchers might consider treated areas of different sizes and apply synthetic control methods separately for each area. However, this approach ignores the spatial structure of the data, and can lead to efficiency loss in spatial settings.
We propose to deal with these issues by developing a Bayesian spatial matrix completion framework that allows us to predict the missing potential outcomes in the different areas around the intervention point while accounting for the spatial structure of the data. Specifically, the missing time series in the absence of treatment for the treated areas are imputed using a weighted average of control time series, where the weights are assumed to vary smoothly over space according to a Gaussian process.
Our motivating application is the construction of the first line of the Florentine tramway, which could have had an effect on the prevalence of businesses in the neighbourhood of the construction site, and at various distances from the tramway stops.

Quantifying the impact of Covid-19 on the Italian labour market

Antonio Pietrafesa[(1)].

(1) University of Rome Tor Vergata, Italy

Abstract: Causal inference in time series has the aim to evaluate the impact of shocks and policy interventions over time. The objective of the present work is to assess the impact of COVID-19 on the Italian labour market by estimating its effect on the time series of activations and terminations of subordinate employment relationships and on the different sectors and types of workers, including fragile categories such as part-time workers, seasonal workers and immigrants. In order to do so, we exploit the information contained in the dataset on compulsory communications (CO) and employ the Bayesian state-space model proposed by Brodersen et al. (2015) with the difference that we consider an exogenous shock in the place of a market intervention. We find that the pandemic shock has had a strong negative effect on activations and terminations with a greater impact on terminations, independently from the disaggregation we consider. We find also a certain degree of heterogeneity in the impacts among genders, sectors, workers and geographical areas and some evidence of a gender gap. On the contrary, we find no significant effects on seasonal workers.

Keywords: Causal inference; Time series; Bayesian; Covid-19.

Non-Linear interaction models: an approach based on varying coefficients

Davide Fabbrico[(1)], Matteo Pedone[(1)], Francesco Claudio Stingo [(1)].

(1) University of Florence, Italy

Abstract: The linear regression model is widely used across various fields of statistics. Numerous extensions have been proposed to enhance its flexibility and broaden its range of applications. In this work, we comprehensively review and study extensions of the linear model in three distinct directions. Firstly, we explore the incorporation of non-linear effects. Secondly, we examine the imposition of a hierarchical structure a priori to select relevant variables. Finally, we aim to enhance the flexibility of the regression coefficients by introducing their dependence on covariates. To evaluate the performance of the linear model under the proposed assumptions, we conduct simulation studies employing four primary metrics, Mean square error (MSE), True positive rate (TPR), False positive rate (FPR) and Matthew's correlation coefficient (MCC). Our goal is to determine whether these assumptions contribute to the improvement of the linear model proposed in the literature.

Keywords: Bayesian Variable Selection; Non-Linear models; Varying Coefficient; Interaction models; Permutation Invariance; Spline Regression.

Bayesian Nonparametrics for Principal Stratification: an Application on Environmental Policies Effects on Health

Dafne Zorzetto[(1)], Falco J. Bargagli-Stoffi[(2)], Antonio Canale [(1)], Fabrizia Mealli [(3)], Francesca Dominici[(2)].

(1) Department of Statistics, University of Padova, Italy,
(2) Department of Biostatistics, Harvard T.H. Chan School of Public Health, Massachusetts, USA,
(3) Department of Economics, European University Institute, Italy

Abstract: Regulatory actions have been enacted in the United States to diminish the levels of pollutants in the air and reduce the connected environmental risks for health. Indirect accountability studies—assessing the causal effect of exposure to higher levels of air pollution—and direct accountability studies—assessing the causal impact of interventions aimed at reducing the level of air pollution—have found solid evidence of health benefits. However, the existing literature lacks robust methods that consider two crucial points in health studies: evaluate heterogeneity in the health effects of air pollution regulations across different groups of individuals, and consider the joint relations between direct and indirect effects. In this work, we develop a novel approach combining Bayesian nonparametric (BNP) methods and Principal stratification (PS) framework to deal with post-treatment variables that are potentially affected by the treatment and also affecting the response. We introduce three major innovations: (i) we rely on BNP methodologies for the imputation of missing potential outcomes for the post-treatment; (ii) we introduce new conditional estimands; (iii) we propose a data-driven methodology to discover causal heterogeneity. We illustrate the performance of the method through simulations. In the application we discover and estimate the heterogeneous effects of US national air quality regulations on pollution levels and health outcomes.

Keywords: Causal Inference; Principal Stratification; Bayesian Nonparametrics; Air Quality Regulations.

Robust inference of causality in high-dimensional dynamical processes from the Information Imbalance of distance ranks

Vittorio Del Tatto [(1)], Gianfranco Fortunato[(1)], Domenica Bueti [(1)], Alessandro Laio[(1,2)].

(1) Scuola Internazionale Superiore di Studi Avanzati, SISSA, via Bonomea 265, 34136 Trieste, Italy,
(2) ICTP, Strada Costiera 11, 34151 Trieste, Italy

Abstract: We introduce an approach which allows inferring causal relationships between variables for which the time evolution is available. Our method builds on the ideas of Granger Causality and Transfer Entropy, but alleviates most of their limitations by being model-free and nonparametric. Specifically, our approach tests whether the predictability of a putative driven system Y can be improved by incorporating information from a potential driver system X, without making assumptions on the underlying dynamics and without the need to compute probability densities of the dynamic variables. Causality is assessed by a rigorous variational scheme based on the Information Imbalance of distance ranks, a recently developed statistical test capable of inferring the relative information content of different distance measures. This framework makes causality detection possible even for high-dimensional systems where only few of the variables are known or measured. Benchmark tests on coupled dynamical systems demonstrate that our approach outperforms other model-free causality detection methods, successfully handling both unidirectional and bidirectional couplings, and it is capable of detecting the arrow of time when present. We also show that the method can be used to robustly detect causality in electroencephalography data in humans.

Keywords: Causality; Time series; Ranks; Nonparametric.

Bayesian estimation of P3b amplitude as a correlate of an internal Signal Detection Theory mechanism

Erik Skjoldan Mortensen [(1)].

(1) University of Copenhagen, Department of Psychology, Denmark

Abstract:
Background: It has previously been shown that the amplitude of the event-related potential P3b [1] correlate well with subjective reports of graded visual perception in response to masked visual stimuli, which has led the P3b to be a candidate neural correlate of consciousness [2]–[4]. Subjective reports of experience have often been collected on ordinal scales, such the Perceptual Awareness Scale (PAS) [5]. Signal Detection Theory has been used to describe how subjects would use such a rating scale in a stimulus presentation/omission discrimination task [6]. The theory describes how an observer would discriminate between the presence and absence of a stimulus, given a low signal-to-noise ratio, by assuming that multiple observations of both stimulus presentation and omission trials generate overlapping distributions of observed values, with any offset in the means of these two distributions indicating above-chance performance of the observer. This theory is employed to simultaneously describe the distributions of observed P3b amplitudes and PAS ratings.
Methods: In a visual task, participants were asked to locate a visually masked stimulus, only shown on a subset of trials, and were asked to rate their experience of seeing the stimulus on a 4-point PAS, ranging from "No experience of seeing a stimulus" to "A clear experience of seeing a stimulus". A total of 16 participants were included, each completing 288 trials of which 2/3 were stimulus presentation trials and 1/3 were stimulus omission trials, with electroencephalography (EEG) captured throughout. Stimuli were displayed for ~33 ms and their intensity were individually adjusted to be just-above threshold for visual perception, with the goal of ensuring participants discriminatory ability being better than random while still producing a large number of both true and false positives, i.e., high PAS ratings during stimulus presentation/omission respectively, and vice versa for both true and false negatives.
Results: The results of a Bayesian multi-level multi-variate regression, modelling P3b amplitude (robust linear) and PAS rating (cumulative probit) as described by trial type (stimulus presentation/omission), show that higher P3b amplitudes are associated with stimulus presentation trials, compared to stimulus omission, and that participants can discriminate between trial types at above-chance level. Analysis of the correlation between the subject-level effects shows a tendency for participants who have better discriminatory ability to also have larger mean differences in P3b amplitude by trial type, as would be expected if SDT describe the observed distribution of P3b amplitudes, as well as the PAS ratings. Two secondary ordinal Bayesian models show that P3b amplitudes has a similar effect on predicting PAS rating during both stimulus presentation and stimulus omission, as would likewise be expected if the proportionate use of each category is described by SDT.
Conclusion: The observed distributions of both PAS ratings and P3b amplitudes are in line with the assumptions of SDT. This work thus further expands upon previous findings that have shown correlations between P3b amplitude and perceptual ratings, by providing a plausible model based on SDT for how the distribution of P3b amplitudes are generated and how they are correlated with PAS ratings.

Keywords: Neural correlate of consciousness, Signal Detection Theory, P3b, EEG.

References

[1] J. Polich, "Updating P300: An integrative theory of P3a and P3b," *Clin Neurophysiol*, vol. 118, no. 10, pp. 2128–2148, 2007, doi: 10.1016/j.clinph.2007.04.019.

[2] C. Dembski, C. Koch, and M. Pitts, "Perceptual awareness negativity: a physiological correlate of sensory consciousness," *Trends Cogn Sci*, vol. 25, no. 8, pp. 660–670, 2021, doi: 10.1016/j.tics.2021.05.009.

[3] J. Förster, M. Koivisto, and A. Revonsuo, "ERP and MEG correlates of visual consciousness: The second decade," *Conscious Cogn*, vol. 80, p. 102917, 2020, doi: 10.1016/j.concog.2020.102917.

[4] S. Dehaene and J.-P. Changeux, "Experimental and Theoretical Approaches to Conscious Processing," *Neuron*, vol. 70, no. 2, pp. 200–227, 2011, doi: 10.1016/j.neuron.2011.03.018.

[5] T. Z. Ramsøy and M. Overgaard, "Introspection and subliminal perception," *Phenomenol Cogn Sci*, vol. 3, no. 1, pp. 1–23, 2004, doi: 10.1023/b:phen.0000041900.30172.e8.

[6] H. Stanislaw and N. Todorov, "Calculation of signal detection theory measures," *Behav Res Methods Instruments Comput*, vol. 31, no. 1, pp. 137–149, 1999, doi: 10.3758/bf03207704.

Median-based Splitting Rules For The Causal Tree

Karolina Gliszczynska[(1)], *Lennard Maßmann*[(2)].

(1) University Duisburg-Essen, Germany,
(2) Ruhr Graduate School in Economics, Germany

Abstract: Our paper contributes to the literature on tree-based methods for causal inference and treatment effect estimation in high-dimensional data. Tree-based methods, like causal trees can be estimated in an honest way, i.e., that model complexity is unrestricted. The imposition of model structure and the estimation referring to a specific model structure is achieved by separate data sets through sample splitting. We present three median-based splitting rules for the causal tree, based on the Median Absolute Deviation (MAD), Least Median Square (LMS) and Median Squared Deviation (MSD), where the Hodges-Lehmann estimator is used either for population median or treatment effect estimation depending on the chosen implementation of the splitting rule. We compare our method with existing ones (mean-based causal tree, transformed outcome tree, fit-based tree, squared t-statistic tree) via a simulation study. Thereby, we draw data from different scenarios, where we distinguish between homo- and heteroscedastic errors as well as correlated and uncorrelated covariates uniformly or standard normally distributed. Our focus is on estimating conditional average treatment effects (CATE) while also considering the model's performance. Moreover, we seek to create reliable confidence intervals for the treatment effects estimated by these variants, without overlooking their accuracy and precision. While our approach performs as good as the existing ones, we receive better confidence interval coverage, so there might be an incentive to deviate from the common used splitting rules based on minimising the mean squared error.

Keywords: causal forest; conditional average treatment effect; splitting rules; conformal inference.

BART as a Gaussian Process

Giacomo Petrillo [(1)].

(1) University of Florence, Italy

Abstract: BART is a nonparametric Bayesian regression technique. The ACIC data challenge has confirmed it as one of the best methods in the context of causal inference for observational studies (Dorie et al. 2019, Thal and Finucane 2023). In a certain limit, BART becomes equivalent to Bayesian Gaussian process (GP) regression. I exploit this connection, both theoretically and practically.
Theoretically, the GP limit offers a different perspective to interpret the nonparametric prior. The GP prior is specified as a covariance function, which is in some ways more legible than the specification as a branching process of BART.
Practically, I show on benchmark datasets that the GP limit surpasses standard BART both on computational time and statistical performance. The GP formulation also features an analytical posterior and easier model building (e.g., for constructions like the Bayesian Causal Forest of Hanh et al. 2020). The main disadvantage is that GP inference scales poorly with dataset size; I think it is possible to overcome this limitation with existing techniques, but I have not made an attempt yet: the software package I provide can realistically be used only up to n = 10 000.
With fixed hyperparameters, on 50% of the 42 benchmark datasets of Chipman et al. (2010), the test set prediction RMSE is the same for standard BART and its GP limit. When different, standard BART is more often at an advantage. However, tuning the hyperparameters, the GP limit becomes faster and more accurate than standard BART, for two reasons: 1) part of the hyperparameter space is precluded to standard BART due to convergence and computational issues in the MCMC, and 2) standard BART requires cross validation to tune the hyperparameters, while the GP limit allows direct inference. The GP limit has been studied before, but the existing literature suggests that it is not profitable in practice (Linero 2017, see also the discussion in Hanh et al. 2020). I hypothesize that my unexpected success is due to an improved calculation of the BART prior covariance function, which requires nontrivial mathematical and computational treatment.

Keywords: BART; Gaussian process; Response surface; Methodology.

Estimating Treatment Effects Using BART Models Within The Conformal Inference Framework

Lennard Maßmann [(1)].

(1) University of Duisburg-Essen / Ruhr Graduate School in Economics, Germany

Abstract: The Bayesian Additive Regression Tree (BART) model as a flexible approach for prediction tasks has shown to perform reasonably well in estimating the Conditional Average Treatment Effect (CATE) and the Individual Treatment Effect (ITE). For instance, thorough empirical evaluations of various BART versions are documented within the Data Analysis Challenge associated with the Atlantic Causal Inference Conference focusing on CATE inference of point estimates when there is significant confounding due to targeted selection. Although prediction intervals of Bayesian models are recovered easily from posterior simulations, sensible uncertainty quantification with good frequentist coverage properties for treatment effects can be challenging. Recently, conformal inference has been introduced to generate trustworthy intervals for uncertainty quantification around ITEs and CATEs within the potential outcome framework. In this work, we review popular BART models (e.g. SBART, XBART, BART-BMA) within the conformal inference framework to produce interval estimates around treatment effects for comparison of coverage properties with uncertainty measures produced by posterior simulations. Numerical studies show coverage issues of traditional machine learning approaches for treatment effect estimation and support the usage of conformal inference for treatment effect uncertainty quantification in various settings.

Keywords: Conformal Inference; BART; Treatment Effect; Uncertainty Quantification.

An investigation of the causal association between social participation and health outcomes: findings from a nationally representative sample using marginal structural models

Shunqi Zhang[(1)], *Nan Zhang*[(1)], *Mark Brown*[(1)].

(1) Department of Social Statistics, The University of Manchester, United Kingdom

Abstract: To assess the effect of different types of social participation on mental health and cognitive function in later life. We used data from four waves of the China Health and Retirement Longitudinal Study (CHARLS) over 7 years (2011-2018). The outcomes were the Centre for Epidemiologic Studies Depression Scale (CES-D 10) and the global cognitive function score. Social participation was measured by eight binary measures. Time-invariant confounders included age, gender, education, and parental educational backgrounds. Time-varying confounders included self-rated health, the score of activities in daily life, the score of mobility, marital status, residential areas, weekly contact with children and the amount of alcohol consumption. We used marginal structural models (MSMs) with stabilised Inverse Probability of Treatment Weights (IPTWs). To deal with the remaining residual imbalance, time-invariant covariates were included in the MSMs. There was evidence of associations between two activities of informal social participation (interaction with friends and caring for a sick adult) and a lower risk of depressive symptoms after adjusting for the confounding. Most of social participation was significantly associated with better cognitive function across waves. Although the possibility of reverse causation cannot be fully excluded, the study has applied MSM with IPTW to address the directionality of the relationship between social participation, cognitive function, and depression..

Keywords: Social participation; Longitudinal analysis; Ageing; Health.