# MACHINE LEARNING FOR CAUSAL INFERENCE

**Period**: a.y. 2026/2027 II semester

**Class times**: Monday - Thursday

**Instructor:**

Prof. Marton Varga
Dept. of Marketing. - Room 4-D1-11
marton.varga@unibocconi.it

## Course description

This PhD course introduces foundational and advanced tools in causal inference and machine learning, with an emphasis on their combined use in empirical research. These tools are increasingly central to high-quality research in management, marketing, and related fields. The course is organized into six sessions. Session 1 reviews core treatment-estimation designs. Session 2 covers prediction, regularization, and applied machine learning. Session 3 examines prediction, policy learning, and the role of machine learning in empirical research. Session 4 turns to causal machine learning and modern treatment-estimation tools, including causal forests, synthetic controls, and double/debiased machine learning. Sessions 5 and 6 focus on modern difference-in-differences, synthetic and stacked difference-in-differences, and empirical applications in management and marketing. The course combines lectures, discussion of recent research papers, and student presentations. As part of class participation, students prepare a synthetic-data analysis implementing one method studied in class, using R, Python, Stata, or another approved language. Participants are expected to actively engage with the material, complete the synthetic-data implementation exercise, present assigned papers, and develop a concept paper. The course is especially well-suited for PhD students looking to apply rigorous empirical methods in their own research projects.

## Course Material

### Session 1: Treatment Estimation

- Abadie, A., & Cattaneo, M. D. (2018). Econometric methods for program evaluation. Annual Review of Economics, 10, 465-503.
- Acemoglu, D., Johnson, S., & Robinson, J. A. (2001). The colonial origins of comparative development: An empirical investigation. American

Economic Review, 91(5), 1369-1401.

- Card, D., & Krueger, A. B. (1994). Minimum wages and employment: Case study of the fast-food industry in New Jersey and Pennsylvania. American Economic Review, 84(4), 772-793.
- Anderson, M., & Magruder, J. (2012). Learning from the crowd: Regression discontinuity estimates of the effects of an online review database. The Economic Journal, 122(563), 957-989.
- Jin, G. Z., & Leslie, P. (2003). The effect of information on product quality: Evidence from restaurant hygiene grade cards. Quarterly Journal of Economics, 118(2), 409-451.

### Session 2: Prediction, Regularization, and Machine Learning

- Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. Journal of Economic Perspectives, 31(2), 87-106.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society Series B, 58(1), 267-288.
- Cutler, A., Cutler, D. R., & Stevens, J. R. (2012). Random forests. In C. Zhang & Y. Ma (Eds.), Ensemble machine learning: Methods and applications (pp. 157-175). Springer.
- Larivière, B., & Van den Poel, D. (2005). Predicting customer retention and profitability by using random forests and regression forests techniques. Expert Systems with Applications, 29(2), 472-484.
- Albuquerque, P., Tusche, A., Varga, M., Gier-Reinartz, N. R., Weber, B., & Plassmann, H. (2026). Do fMRI data improve predictions of product adoption by store managers and sales per store of consumer packaged goods? Journal of Marketing Research. Advance online publication.

### Session 3: Machine Learning, Policy Learning, and Causal Inference

- Athey, S. (2019). The impact of machine learning on economics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), The economics of artificial intelligence: An agenda (pp. 507-547). University of Chicago Press.
- Athey, S. (2017). Beyond prediction: Using big data for policy problems. Science, 355(6324), 483-485.
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Obermeyer, Z. (2015). Prediction policy problems. American Economic Review Papers and Proceedings, 105(5), 491-495.
- Varian, H. R. (2014). Big data: New tricks for econometrics. Journal of Economic Perspectives, 28(2), 3-28.
- Athey, S., Imbens, G. W., Pham, T., & Wager, S. (2017). Estimating average treatment effects: Supplementary analyses and remaining challenges. American Economic Review Papers and Proceedings, 107(5), 278-281.
- Athey, S., Blei, D., Donnelly, R., Ruiz, F. J., & Schmidt, T. (2018). Estimating heterogeneous consumer preferences for restaurants and travel time using mobile location data. AEA Papers and Proceedings, 108,

64-67.

- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., & Mullainathan, S. (2018). Human decisions and machine predictions. Quarterly Journal of Economics, 133(1), 237-293.

### Session 4: Causal Machine Learning and Modern Treatment Estimation

- Jawadekar, N., Kezios, K., Odden, M. C., Stingone, J. A., Calonico, S., Rudolph, K., & Zeki Al Hazzouri, A. (2023). Practical guide to honest causal forests for identifying heterogeneous treatment effects. American Journal of Epidemiology, 192(7), 1155-1165.
- Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. Journal of the American Statistical Association, 105(490), 493-505.
- Ben-Michael, E., Feller, A., & Rothstein, J. (2022). Synthetic controls with staggered adoption. Journal of the Royal Statistical Society Series B, 84(2), 351-381.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., & Newey, W. (2017). Double/debiased/Neyman machine learning of treatment effects. American Economic Review Papers and Proceedings, 107(5), 261-265.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. The Econometrics Journal, 21(1), C1-C68.
- Ahrens, A., Chernozhukov, V., Hansen, C., Kozbur, D., Schaffer, M., & Wiemann, T. (2026). An introduction to double/debiased machine learning. arXiv:2504.08324.
- Ellickson, P. B., Kar, W., & Reeder, J. C., III. (2023). Estimating marketing component effects: Double machine learning from targeted digital promotions. Marketing Science, 42(4), 704-728.

### Session 5: Modern Difference-in-Differences and Event Studies

- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics, 225(2), 175-199.
- Baker, A. C., Larcker, D. F., & Wang, C. C. Y. (2022). How much should we trust staggered difference-in-differences estimates? Journal of Financial Economics, 144(2), 370-395.
- Li, Z., & Wang, G. (2025). On-demand delivery platforms and restaurant sales. Management Science, 71(7), 5788-5804.
- He, S., Hollenbeck, B., & Proserpio, D. (2022). The market for fake reviews. Marketing Science, 41(5), 896-921.
- Adalja, A., Liaukonytė, J., Wang, E., & Zhu, X. (2023). GMO and non-GMO

labeling effects: Evidence from a quasi-natural experiment. Marketing Science, 42(2), 233-250.

### Session 6: Synthetic and Stacked Difference-in-Differences

- Wing, C., Freedman, S. M., & Hollingsworth, A. (2024). Stacked difference-in-differences. NBER Working Paper No. 32054.
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference-in-differences. American Economic Review, 111(12), 4088-4118.
- Guyt, J. Y., van Lin, A., & Keller, K. O. (2025). Banning unsolicited store flyers: Does helping the environment hurt retailing? Marketing Science, 44(5), 1104-1124.
- Lambrecht, A., Tucker, C., & Zhang, X. (2024). TV advertising and online sales: A case study of intertemporal substitution effects for an online travel platform. Journal of Marketing Research, 61(2), 248-270.
- Li, X., Deng, Y., Manchanda, P., & De Reyck, B. (2025). Can lower(ed) expert opinions lead to better consumer ratings? The case of Michelin stars. Management Science, 72(3), 2427-2450.

## Assessment Methods

### Class Participation (50%)

- Includes attendance, presentation of papers assigned by the instructor, thoughtful engagement with assigned readings, active contribution to class discussion, and the practical implementation exercise.
- Practical Implementation Exercise – As part of class participation, choose one method studied in class, write reusable code in R, Python, Stata, or another approved language, generate a synthetic dataset, and conduct a short analysis using that method. The topic and method for the synthetic-data exercise are chosen by the student after agreement with the instructor; they do not need to match the topic or method of the concept paper. The analysis is to be presented in class by the student.

### Concept Paper (50%)

A research proposal (max 10 pages) on an empirical topic in management or marketing using observational data and a method covered in class. Emphasis is placed on empirical design and methodological clarity.

### Paper Requirements:

- Literature Contribution – Define your research gap.
- Data Description – Specify observational dataset(s), data structure, and unit of observation.
- Empirical Approach – Describe methodology with equations/visuals.
- Expected Results – Briefly describe anticipated findings.

### Use of Generative AI

Generative AI tools may be used to help illustrate the chosen method on student-generated synthetic data and to support coding for the Practical Implementation Exercise. Students remain responsible for the content, code, analysis, and outcome.

For paper presentations, Generative AI use is limited to polishing slide tables, text, and pictures. Students must present the papers assigned by the instructor and remain responsible for the content.

For the concept paper, Generative AI may be used only to improve English; it is not allowed for any other purpose.