

Placement Director: Professor Alessandro Pavan 847-491-8266 alepavan@northwestern.edu
 Placement Administrator: Lola May Ittner 847-491-5213 econjobmarket@northwestern.edu

Contact Information

Department of Economics
 Northwestern University
 2211 Campus Drive
 Evanston, IL 60208

Mobile: 773-290-9341
amilcare@u.northwestern.edu
<https://sites.northwestern.edu/amilcarvelez/>
 Citizenship: Perú

Fields

Econometrics

Education

Ph.D., Economics, Northwestern University (anticipated) 2025
 Committee: Ivan A. Canay (co-chair), Federico A. Bugni (co-chair), Joel L. Horowitz
 M.A., Economics, Northwestern University 2020
 M.A., Mathematics, IMPA, Brazil 2015
 B.A., Economics, PUCP, Perú 2014

Fellowships & Awards

Dissertation University Fellowship, Northwestern University 2024–2025
 Robert Eisner Graduate Fellowship, Northwestern University 2023–2024
 University Fellowship by The Graduate School, Northwestern University 2019–2020
 National Research Council Master Student Fellowship, IMPA 2015
 Solidary Academic Excellence Scholarship (BEAS), PUC-Perú 2011–2014
 Bronze Medal in the 49th International Mathematical Olympiad, Spain 2008

Teaching Experience

Lecturer, Northwestern University
 Econometrics Review for Incoming Ph.D. Students 2023Su
 Teaching Assistant, Northwestern University
 Econ 281: Introduction to Econometrics (Undergraduate) 2022W, 2023Sp
 Econ 329: Experimental Economics (Undergraduate) 2022F
 Econ 480-3: Econometrics (Graduate) 2021Sp, 2022Sp
 Econ 480-1: Econometrics (Graduate) 2021F

Research Experience

Research Assistant, Professor Federico A. Bugni, Northwestern University 2022Su, 2023W
 Research Assistant, Professor Ivan A. Canay, Northwestern University 2020F, 2021W

Conferences & Presentations

2024: Notre Dame

Refereeing

AER: Insights, Journal of Business & Economic Statistics

Job Market Paper

“On the Asymptotic Properties of Debiased Machine Learning Estimators”

This paper studies the properties of debiased machine learning (DML) estimators under a novel asymptotic framework, offering insights for improving estimator performance in applications. DML is an estimation method suited to economic models in which the parameter of interest depends on unknown nuisance functions that must be estimated. It requires weaker conditions than previous methods while still ensuring standard asymptotic properties. Existing theoretical results do not distinguish between the two alternative versions of DML estimators, namely, DML1 and DML2. Under a new asymptotic framework, this paper demonstrates that DML2 asymptotically dominates DML1 in terms of bias and mean squared error, formalizing a previous conjecture based on the simulation results of their relative performance. Additionally, this paper provides guidance for improving DML2 performance in applications.

Publications

“A User’s Guide for Inference in Models Defined by Moment Inequalities” with Ivan A. Canay and Gaston Illanes

Accepted at the *Journal of Econometrics*, 2024+

Models defined by moment inequalities have become a standard modeling framework for empirical economists, spreading over a wide range of fields within economics. From the point of view of an empirical researcher, the literature on inference in moment inequality models is large and complex, including multiple survey papers that document the non-standard features these models possess, the main novel concepts behind inference in these models, and the most recent developments that bring advances in accuracy and computational tractability. In this paper we present a guide to empirical practice intended to help applied researchers navigate all the decisions required to frame a model as a moment inequality model and then to construct confidence intervals for the parameters of interest. We divide our template into four main steps: (a) a behavioral decision model, (b) moving from the decision model to a moment inequality model, (c) choosing a test statistic and critical value, and (d) accounting for computational challenges. We split each of these steps into a discussion of the “how” and the “why”, and then illustrate how to take these steps to practice in an empirical application that studies identification of expected sunk costs of offering a product in a market.

“The A/B Testing Problem with Gaussian Priors” with Eduardo M. Azevedo, David Mao, and José Luis Montiel Olea

Journal of Economic Theory, 2023, Vol. 210, 1–23

A risk-neutral firm can perform a randomized experiment (A/B test) to learn about the effects of implementing an idea of unknown quality. The firm’s goal is to decide the experiment’s sample size and whether or not the idea should be implemented after observing the experiment’s outcome. We show that when the distribution for idea quality is Gaussian and there are linear costs of experimentation, there are exact formulae for the firm’s optimal implementation decisions, the value of obtaining more data, and optimal experiment sizes. Our formulae—which assume that companies use randomized experiments to help them maximize expected profits—provide a simple alternative to i) the standard rules-of-thumb of power calculations for determining the sample size of an experiment, and also to ii) ad hoc thresholds based on statistical significance to interpret the outcome of an experiment.

“On the Robustness to Misspecification of α -posteriors and their Variational Approximations” with Marco Avellana, José Luis Montiel Olea, and Cynthia Rush

Journal of Machine Learning Research, 2022, Vol. 23 (147), 1–51

α -posteriors and their variational approximations distort standard posterior inference by downweighting the likelihood and introducing variational approximation errors. We show that such distortions, if tuned appropriately, reduce the Kullback–Leibler (KL) divergence from the true, but perhaps infeasible, posterior distribution when there is potential parametric model misspecification. To make this point, we derive a Bernstein–von Mises theorem showing convergence in total variation distance of α -posteriors and their variational approximations to limiting Gaussian distributions. We use these limiting distributions to evaluate the KL divergence between true and reported posteriors. We show that the KL divergence is minimized by choosing α strictly smaller than one, assuming there is a vanishingly small probability of model misspecification. The optimized value of α becomes smaller as the misspecification becomes more severe. The optimized KL divergence increases logarithmically in the magnitude of misspecification and not linearly as with the usual posterior. Moreover, the optimized variational approximations of α -posteriors can induce additional robustness to model misspecification beyond that obtained by optimally downweighting the likelihood.

“Posterior Distribution of Non-differentiable Functions” with Toru Kitagawa, José Luis Montiel Olea, and Jonathan Payne

Journal of Econometrics, 2020, Vol. 217 (1), 161–175

This paper examines the asymptotic behavior of the posterior distribution of a possibly nondifferentiable function $g(\theta)$, where θ is a finite-dimensional parameter of either a parametric or semiparametric model. The main assumption is that the distribution of a

suitable estimator $\widehat{\theta}_n$, its bootstrap approximation, and the Bayesian posterior for θ all agree asymptotically. It is shown that whenever g is locally Lipschitz, though not necessarily differentiable, the posterior distribution of $g(\theta)$ and the bootstrap distribution of $g(\widehat{\theta}_n)$ coincide asymptotically. One implication is that Bayesians can interpret bootstrap inference for $g(\theta)$ as approximately valid posterior inference in a large sample. Another implication—built on known results about bootstrap inconsistency—is that credible intervals for a nondifferentiable parameter $g(\theta)$ cannot be presumed to be approximately valid confidence intervals (even when this relation holds true for θ).

Other papers

“The Local Projection Residual Bootstrap for AR(1) Models”

Revision requested by *Econometric Theory*

This paper proposes a local projection residual bootstrap method to construct confidence intervals for impulse response coefficients of AR(1) models. Our bootstrap method is based on the local projection (LP) approach and involves a residual bootstrap procedure applied to AR(1) models. We present theoretical results for our bootstrap method and proposed confidence intervals. First, we prove the uniform consistency of the LP-residual bootstrap over a large class of AR(1) models that allow for a unit root. Then, we prove the asymptotic validity of our confidence intervals over the same class of AR(1) models. Finally, we show that the LP-residual bootstrap provides asymptotic refinements for confidence intervals on a restricted class of AR(1) models relative to those required for the uniform consistency of our bootstrap.

“The out-of-sample prediction error of the \sqrt{LASSO} and related estimators” with José Luis Montiel Olea, Cynthia Rush, and Johannes Wiesel

Under Major Revision at *Annals of Statistics*

We study the classical problem of predicting an outcome variable, Y , using a linear combination of a d -dimensional covariate vector, \mathbf{X} . We are interested in linear predictors whose coefficients solve:

$$\inf_{\beta \in \mathbb{R}^d} \left(\mathbb{E}_{\mathbb{P}_n} [|Y - \mathbf{X}^\top \beta|^r] \right)^{1/r} + \delta \rho(\beta),$$

where $\delta > 0$ is a regularization parameter, $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_+$ is a convex penalty function, \mathbb{P}_n is the empirical distribution of the data, and $r \geq 1$. Our main contribution is a new bound on the out-of-sample prediction error of such estimators.

The new bound is obtained by combining three new sets of results. First, we provide conditions under which linear predictors based on these estimators solve a *distributionally robust optimization* problem: they minimize the worst-case prediction error over distributions that are close to each other in a type of *max-sliced Wasserstein metric*. Second, we provide a detailed finite-sample and asymptotic analysis of the statistical properties of the balls of distributions over which the worst-case prediction error is analyzed. Third, we present an oracle recommendation for the choice of regularization parameter, δ , that guarantees good out-of-sample prediction error.

“Identification and Inference on Treatment Effects under Covariate-Adaptive Randomization and Imperfect Compliance” with Federico A. Bugni, Mengsi Gao, and Filip Obradović

Submitted

Randomized controlled trials (RCTs) frequently use covariate-adaptive randomization (CAR) (e.g., stratified block randomization) and commonly suffer from imperfect compliance. This paper studies the identification and inference for the average treatment effect (ATE) and the average treatment effect on the treated (ATT) in such RCTs with a binary treatment.

We first develop characterizations of the identified sets for both estimands. Since data are generally not i.i.d. under CAR, these characterizations do not follow from existing results. We then provide consistent estimators of the identified sets and asymptotically valid confidence intervals for the parameters. Our asymptotic analysis leads to concrete practical recommendations regarding how to estimate the treatment assignment probabilities that enter in estimated bounds. In the case of the ATE, using sample analog assignment frequencies is more efficient than using the true assignment probabilities. On the contrary, using the true assignment probabilities is preferable for the ATT.

“On the Power Properties of Confidence Sets for Parameters with Interval Identified Sets” with Federico A. Bugni, Mengsi Gao, and Filip Obradović

Submitted

This paper studies the power properties of confidence intervals (CIs) for a partially-identified parameter of interest with an interval identified set. We assume the researcher has bounds estimators to construct the CIs proposed by Stoye (2009), referred to as CI_α^1 , CI_α^2 , and CI_α^3 . We also assume that these estimators are “ordered”: the lower bound estimator is less than or equal to the upper bound estimator.

Under these conditions, we establish two results. First, we show that CI_α^1 and CI_α^2 are equally powerful, and both dominate CI_α^3 . Second, we consider a favorable situation in which there are two possible bounds estimators to construct these CIs, and one is more efficient than the other. One would expect that the more efficient bounds estimator yields more powerful inference. We prove this desirable result holds for CI_α^1 and CI_α^2 , but not necessarily for CI_α^3 .

“Linear IV Model with Missing Data on the Instrumental Variable”

This paper studies the problem of identification in an IV model with missing data on the instrumental variable. I consider an agnostic stance on the distribution of the missing data and a worst-case scenario approach to confront the missing data problem. First, I characterize the identified set of the parameter of interest and make explicit that this is an extremely complex object to compute. Next, following the literature on partial identification, I propose an outer identified set—a superset of the identified set that is easier to compute. This outer identified set is based on a moment inequality model. Then, I show that, under some testable assumptions, this outer identified set is equal to the identified set.

“Representation Learning in Linear Factor Models” with José Luis Montiel Olea

In this work, we analyze recent theoretical developments in the representation learning literature through the lens of a linear Gaussian factor model. First, we derive *sufficient representations*—defined as functions of covariates that, upon conditioning, render the outcome variable and covariates independent. Then, we study the theoretical properties of these representations and establish their *asymptotic invariance*; which means the dependence of the representations on the factors’ measurement error vanishes as the dimension of the covariates goes to infinity. Finally, we use a decision-theoretic approach to understand the extent to which representations are useful for solving *downstream tasks*. We show that the conditional mean of the outcome variable given covariates is an asymptotically invariant and sufficient representation that can solve *any task* efficiently, not only prediction.

Professional Experience

Central Bank of Perú (BCRP)
Macro Modelling Division
Financial Stability Division
(on leave 2015)

May 2018 - Jul 2019
May 2014 - Apr 2018

Languages

English (fluent), Spanish (native)

References

Professor Ivan A. Canay
Department of Economics
Northwestern University
2211 Campus Drive
Evanston, IL 60208
847.491.2929
iacanay@northwestern.edu

Professor Federico A. Bugni
Department of Economics
Northwestern University
2211 Campus Drive
Evanston, IL 60208
847.491.8212
federico.bugni@northwestern.edu

Professor Joel L. Horowitz
Department of Economics
Northwestern University
2211 Campus Drive
Evanston, IL 60208
847.491.8253
joel-horowitz@northwestern.edu