Rafael Pintro Schmitt

Application to the NBER Digital Economics and Artificial Intelligence Tutorial

Email	rafaelpintroschmitt@berkeley.edu
Institution	University of California, Berkeley, Department of Economics
Graduate study	Ph.D. in Economics, August 2024–present (third year)
Research adviser	Patrick Kline
Dissertation committee	Not yet finalized
Research interests	Economics of AI; measurement, prediction, and fairness

Graduate Study and Coursework

I have completed two years of coursework in my ongoing Economics PhD, including the graduate core in economic theory, macroeconomics, and econometrics, as well as field courses in labor and IO. My cumulative graduate GPA is **3.975/4** (61 completed units). I received the 2026 E. J. Swan Prize from the Berkeley Economics Department.

Term	Course	Title	Units	Grade
Fall 2024	ECON 201A	Economic Theory	4	A
	ECON 202A	Macroeconomics	4	A
	ECON 240A	Econometrics	5	A
	ECON 295	Survey of Research in Economics	1	S
Spring 2025	ECON 201B	Economic Theory	4	A
	ECON 202B	Macroeconomics	4	A–
	ECON 210A	Introduction to Economic History	3	A
	ECON 240B	Econometrics	4	A
	ECON 295	Survey of Research in Economics	1	S
Fall 2025	ECON 219A	Foundations of Psychology and Economics	3	A
	ECON 220A	Industrial Organization	3	A+
	ECON 250A	Labor Economics	3	A
	ECON 301	GSI Practicum	4	S
	ECON 375	GSI Pedagogy Workshop	2	S
Spring 2026	ECON 215B	Political Economics	3	A+
	ECON 220B	Industrial Organization	3	A+
	ECON 244	Applied Econometrics	3	A
	ECON 301	GSI Practicum	4	S
	PHDBA 279PA	Political Economy: Frameworks	3	A+

Courses and grades reproduced from the UC Berkeley academic summary; *S* denotes satisfactory. Only completed coursework through Spring 2026 is listed.

Research Interests and Ongoing Projects

My research interests lie at the intersection of labor economics and applied econometrics. I am particularly interested in how AI will affect the measurement of economic phenomena and accelerate scientific progress, how it will change the nature of work, and who the winners and losers of technological change are.

1. Economic Effects of Generative AI

I am working with Germán Reyes, Sebastian Otero, Cauê Dobbin, Nano Barahona and Joaquin Serrano on a project selected for OpenAI's Economic Research Exchange, a research program supporting empirical work on the economic effects of AI. We will study how generative AI affects employment in Brazil. Existing evidence on employment effects is mixed and comes almost entirely from high-income countries. Moreover, most studies rely on occupational exposure - what AI could do - rather than observed usage - whether and how workers actually use it.

The project combines RAIS, Brazil's matched employer-employee census of formal employment; PNAD Contínua, a national household survey covering both formal and informal work; and ChatGPT usage data for Brazil accessed through a partnership with OpenAI's Office of the Chief Economist. Together, these sources allow us to study labor-market adjustment beyond formal employment records and to distinguish the economic effects of realized AI use from those associated with predicted occupational exposure.

We ask three questions. First, does AI exposure affect formal-sector hiring, separations, and earnings, and how do these effects compare with evidence from the United States? Second, where do displaced workers go - into informal employment, non-employment, or new formal jobs potentially created by productivity gains? Third, how does observed AI usage, rather than predicted exposure alone, relate to displacement and earnings? More broadly, the project investigates whether the employment consequences of AI differ in economies where informality is an important margin of adjustment.

2. Undisclosed AI Use in Judicial Decision-Making

I am developing a project on generative AI adoption in the Brazilian judiciary, focusing on usage that judges do not explicitly disclose. Courts are a useful setting because AI may increase the productivity of legal work while potentially affecting the reasoning, consistency, or substantive content of decisions. Yet conventional measures of adoption - surveys and self-reports - and RCTs encouraging the rollout of a tool may fail to capture informal or unacknowledged use, which is likely much more prevalent.

The first stage is a measurement problem: can one construct a credible panel measure of AI-assisted judicial writing from case documents? I have scraped the text of over 50 million court dockets from the state of São Paulo in Brazil, and have partnered with Pangram to construct a judge-level panel of AI-assisted judicial writing.

I plan to study who adopts, how adoption diffuses, and what changes after adoption. A particularly interesting question is whether AI standardizes the application of the law. If judges receive similar assistance in interpreting cases and drafting decisions, adoption could reduce differences in the application of the law across judges facing similar cases and attenuate pre-existing biases across defendants or case types. However, AI could reinforce systematic biases in the information or legal reasoning it provides, or allow judges to implement their existing subjective preferences more efficiently.

3. Measurement, Prediction, and the Evaluation of AI

A third part of my research studies how to design and evaluate measurement systems when the object being measured is latent and the available signals may be systematically biased. In an ongoing project using item-level data from Brazil's national secondary-school examination (ENEM), I examine how university-admissions scoring rules can be redesigned to reduce demographic inequality without sacrificing their ability to assess academic preparation.

The setting is particularly useful because a test score is itself a prediction: individual item responses are aggregated into a one-dimensional measure intended to rank students on an underlying attribute. But questions may differ in informativeness and may also exhibit differential performance across demographic groups. The usual strategies for diagnosing bias often require external validation outcomes, some unbiased questions, or strong assumptions about ability distributions. In many applications, none of those anchors is available.

Our approach uses out-of-sample prediction on held-out exam questions to evaluate alternative scoring rules. We study the frontier between predictive quality and demographic dependence, and derive bounds on changes in the preparation of admitted students under an assumption of no average reverse discrimination, without assuming that any particular question is unbiased. The project is motivated by a practical decision problem: how much can a scoring system be altered to improve equality of opportunity before it risks compromising the attribute it is designed to measure?

I see this problem as closely related to the evaluation of AI systems. AI is increasingly used to score, rank, predict, and recommend, but high-stakes applications rarely come with an unquestioned ground truth. Labels may reflect earlier institutional decisions, performance may have multiple dimensions, and conventional accuracy measures can obscure changes in who benefits from a decision rule. These concerns arise in education, hiring, credit, and public administration, as well as in assessing AI-generated work itself.

Why the NBER Tutorial

The NBER tutorial is particularly well aligned with these projects. I would value the opportunity to learn from researchers studying the economics of AI adoption and its organizational effects, and to discuss the design and interpretation of empirical measures of AI usage and performance. Participation in the tutorial and the accompanying Digital Economics and AI meetings would help me connect the empirical questions in my current work to a broader research agenda on the diffusion, consequences, and evaluation of artificial intelligence.