December 1, 2025

FABIAN ROEBEN

Ph.D. Candidate | Department of Economics | University of Cologne

Weyertal 92
50931 Cologne, Germany
f.roeben@wiso.uni-koeln.de
+49 157 38199497

Personal Website
GitHub
Google Scholar
LinkedIn

Application: Digital Economics and AI Tutorial

Dear NBER Digital Economics and AI Tutorial Organizers,

I am writing to apply for the NBER Digital Economics and AI Tutorial at Stanford University. The tutorial's focus on digital economics, digital platforms, and artificial intelligence closely aligns with my research on (i) the impact of AI via custom experimental platforms and (ii) AI, machine learning, and natural language processing for the economic measurement of latent structures.

Graduate studies. I am a Ph.D. candidate in Economics at the University of Cologne since 2020 and am currently in the research stage of the program. I have completed two years of Ph.D. coursework, including advanced classes in machine learning, econometrics, experimental methods, and digital economics. I expect to complete my dissertation in 2026. Prior to the Ph.D., I obtained an M.A. in Philosophy from University College London and an M.Sc. in Finance from the Stockholm School of Economics.

Research interests. My research lies at the intersection of applied AI, media, and behavioral economics and has two main strands. First, I design and deploy behaviorally informed AI applications, especially mobile apps, as platforms in large-scale field experiments to study the impacts of AI, for example in media markets. Second, I use AI, machine learning, and natural language processing to measure latent structures, for example in behavioral traits, expectations, and innovation from rich text, survey, and administrative data. Across those strands, my work treats AI both as a technology that reshapes markets and as a tool for economic discovery.

Research projects. In the first strand, "News Customization with AI" (with Chopra, Haaland, Roth, and Sticher) asks how AI-enabled customization of news presentation affects what people choose to read and how satisfied they are with their news diet. I built a mobile news app that ingests Associated Press articles and allows users to tailor the presentation of each article, along dimensions such as complexity, sentiment, extent of opinion, entertainment, and political perspective, while holding the underlying news event fixed, using a multi-agent AI rewriting and validation system. In large-scale field experiments with several thousand U.S. news consumers, customization uncovers substantial heterogeneity in news preferences and enables users to correct mismatches between their current news supply and desired style. Relative to a control group that reads a fixed version of each article, users offered customization converge to more fact-driven and less opinionated content and report significantly higher news satisfaction, with comparable overall usage. The project demonstrates how generative AI can be used not only to create content, but also to unbundle presentation from coverage and reduce matching frictions in news markets.

In related work, "The Asymmetric Response to Political Spin" (with Chopra, Haaland, Roth, and Sticher), we study how political spin—the strategic framing of identical facts in liberal, conservative, or neutral language—shapes news consumption and political attitudes. Using the same AI-powered news app and rewriting infrastructure, we generate liberal, conservative, and neutral versions of Associated Press articles while holding underlying events constant and randomly assign users to aligned, misaligned, or neutral spin. This design isolates the causal effect of spin itself, separate from fact selection, and allows us to connect variation in spin to differences in engagement, perceived accuracy, trust, and polarization.

The second strand of my research uses AI, machine learning, and NLP as tools to uncover latent structure in preferences and innovation. "How Far Do Latent Traits Travel?" studies a panel of 2,058 U.S. adults who complete baseline preference, cognition, personality, and decision tasks plus follow-up outcome measures across economic, social, technological, fairness, and political domains. I estimate a factor model that separates stable trait variance from measurement error by exploiting a retest wave and compare the predictive power of the resulting traits to flexible machine-learning models and an LLM-based "digital twin." Trait-based summaries perform similarly to these benchmarks and nearly exhaust predictable variation in many economic and social domains, but leave substantial residual variation in political and evaluative ones, highlighting domains where additional, domain-specific measurement remains essential.

A complementary project, "De-biasing Discovery: Algorithms to Uncover Missed Opportunities", develops an ex-ante system that scores individual research papers, using only information available at publication, on their likelihood and timing of translation into patented technologies. The model is trained on a global AI corpus linking millions of papers and patents. I use the model to define "missing translations": papers the algorithm scores as high-potential that do not translate within a given window, and to map in which fields and institutional categories such gaps arise. Ongoing work uses interpretable analyses of textual and structural features, such as concreteness and application-oriented phrasing, conditional on topic and venue, to understand which cues help high-novelty work translate into patents and to design venue-neutral, calibrated rules that could redirect limited review capacity toward overlooked, high-potential papers.

Additional projects extend this measurement perspective to expectations, expert decisions, and economic behavior. In "Macroeconomic Expectations and Household Types: Evidence from Panel Data", I use expectations data for 11,302 U.S. households to cluster households into types based on the joint structure of their macroeconomic expectations and study how these types shape the response of planned consumption to changes in inflation expectations. "Experts Under Pressure: Information Aggregation and Misallocation in Emergency Triage" analyzes approximately 420,000 emergency department visits to examine whether triage nurses and physicians efficiently aggregate hard metrics (such as vitals and age) and soft information (free-text complaints) when assigning acuity and deciding on admissions. Using machine-learning benchmarks as reference points for possible misallocation, I investigate biases in human decision-making.

Finally, "Self-Revising Language Models for Learning Economic Decision Rules" adapts GEPA, a reflective prompt-optimization algorithm from recent computer science work, to an economic setting. The project treats the language model as an agent and the prompt as its "theory": given an economic scenario, the prompt elicits both a prediction and an explanation in words. The model then repeatedly revises its own prompts to learn a small set of accurate, concise, and diverse language-level decision rules. I apply this framework to economically meaningful domains such as risky choice.

Dissertation. My dissertation, titled "Essays in Digital Economics and AI," is supervised by Prof. Christopher Roth and Prof. Tom Zimmermann and combines the two themes described above: (i) AI-powered digital platforms as experimental laboratories to study the impact of AI, and (ii) AI/ML/NLP as measurement tools to uncover latent structure in traits, expectations, and missed opportunities in innovation. The core chapters are "News Customization with AI," "The Asymmetric Response to Political Spin," "How Far Do Latent Traits Travel?" and "De-biasing Discovery: Algorithms to Uncover Missed Opportunities." I expect to complete the dissertation by the end of 2026.

Thank you for considering my application. I would be delighted to participate in the Digital Economics and AI Tutorial and the associated NBER Digital Economics and AI meeting and look forward to learning from, and contributing to, the discussions there.

Sincerely,
Fabian Roeben