Victor de Chaisemartin: Application for Digital Economics and AI tutorial, Spring 2027

Contact Information

Full name: Victor Yaneng Wang de Chaisemartin
Academic email address: vyw20@mit.edu
Personal email address: victor.dcm57@gmail.com
Phone number: 617-216-8116
Office: Building E62-390, MIT Sloan School of Management

Description of PhD program

I am a third-year PhD student in the Information Technology group at MIT Sloan, with my primary advisor being John Horton. My research sits at the intersection of AI, behavioural economics and econometrics.

I have completed the first two years of coursework required by my program, and have passed my qualifying exams.

My previous coursework includes:

- The full 1st-year PhD microeconomics and econometric sequences offered by the MIT economics department
- PhD machine learning classes with the MIT computer science department
- Digital Economics and Economics of AI: seminar with Catherine Tucker
- Economics of Information Technology: seminar with John Horton and Alessandro Acquisti
- Algorithms and Behavioural Science: taught by Sendhil Mullainathan and Ashesh Rambachan
- Research Design: Decision-Theoretic Perspectives: seminar with Dean Eckles

Description of research interests

I am currently exploring how modern AI capabilities enable new ways of thinking about behavioural economics. I am developing this through two related lines of work. First, I am building AI pipelines that generate novel models and representations of human decision-making. Second, I am exploring new ways of designing experiments with AI that enable us to collect novel behavioural evidence.

Publications

Numbers Tell, Words Sell (with Michael Thaler and Mattie Toma). Forthcoming at the *American Economic Review*.

Working paper link here.

This paper studies what happens when experts have strategic incentives to communicate numeric evidence in natural language (e.g. 'probable') rather than in numbers (e.g. '60%'). In particular, when experts are incentivised to persuade audiences by over- or under-stating the truth, they may leverage the imprecision of language to slant the truth while maintaining plausible deniability. We run two experiments: one where members of the general public communicate probabilities of abstract events, and a second where academic researchers communicate scientific findings to policymakers. We find that incentives to persuade increase the likelihood of experts using language by 25-29 percentage points; furthermore, this successfully persuades audiences. Experts who prefer language exhibit greater slanting of language than numbers.

Working papers

Finding Subgroups with Significant Treatment Effects (with Jann Spiess and Vasilis Syrgkanis)

arXiv link to working paper here.

Social science researchers and policymakers often evaluate interventions using RCTs, but outcomes can be noisy, leading to small or imprecisely estimated average treatment effects. This paper instead proposes new machine-learning tools that provide reliable evidence of an intervention's efficacy by finding *subgroups* with statistically significant treatment effects, even if overall effects are insignificant.

We are currently applying these tools to discover novel evidence on behavioural outcomes from a development economics RCT in Ethiopia (Blattman and Dercon, 2018). Our methods could also be applied to A/B testing on online platforms where there is rich but noisy data.

Work in progress

Inducting new models of social preferences using machine learning (with Ashesh Rambachan)

In this project, we are building an AI pipeline that will induct new models of social preferences from experimental datasets across various standard settings (e.g. ultimatum and dictator games). This builds on Mullainathan and Rambachan (2026), who use machine learning to automatically generate anomalies for existing economic models. This project will push that approach one step further, by using those anomalies as the basis to generate new theories with AI.

Eliciting causal graphs using AI interviewers (with John Horton and Ben Manning)

This project uses John Horton's Expected Parrot platform to recruit subjects and ask them how they would reason causally about a decision problem - e.g. given a set of five job candidates, how would they reason about which one to hire? The survey is conducted automatically by an AI interviewer, which then converts the subjects' responses into a 'causal graph' showing the elicited decision variables and causal relations between them. We are currently prototyping how such causal graphs could be used to identify contradictions in subjects' reasoning and improve their decision-making processes.

Using LLM simulations as Bayesian priors for experimental design

Suppose we have many possible treatment arms that we could potentially include in a full-scale experiment, but due to logistical constraints we can only actually implement a subset of them. For instance, we might be interested in poverty alleviation, and have a large range of possible treatments including cash transfers, microcredit, vaccinations, etc.

To decide which interventions to include, we could first write down a prior over the effectiveness of different treatment arms, then update this using pilot data. However, obtaining a sensible prior in this high-dimensional space is itself difficult.

This project proposes using *LLM agent simulations* of the experiment to derive the prior. Furthermore, uncertainty in the prior belief can be modeled by considering how estimated treatment effects co-vary over a range of feasible LLM simulation implementations.