

YONGHONG ZHANG

Causal Inference · Reliable AI Agents · Climate Finance

PhD Candidate in Economics & Business, Universidad Autónoma de Madrid · Madrid, Spain

yonghong.zhang@estudiante.uam.es | linkedin.com/in/yonghong-zhang | github.com/yonghongzhang-io

PROFILE

PhD candidate working at the intersection of causal inference, climate finance, and reliable AI for empirical research. I study whether AI agents genuinely perform execution-critical research tasks and whether their outputs are evaluated against valid criteria. My current work spans agent-benchmark measurement validity, execution-grounded evaluation of causal workflows, evidence-grounded auditing of DID studies, and multi-agent measurement of firm-level ESG perceptions from multilingual news. My empirical research examines the environmental co-benefits of China's emissions-trading pilots using DID and spatial econometrics.

RESEARCH INTERESTS

Causal inference and policy evaluation; reliable AI and agent evaluation; climate finance and carbon markets; firm-level ESG measurement; AI-assisted financial decision-making; evidence-grounded empirical research.

EDUCATION

PhD in Economics & Business, Universidad Autónoma de Madrid Feb 2024 – Feb 2028 (expected)

Madrid, Spain. Focus: carbon markets, climate finance, environmental policy evaluation, causal inference, AI-assisted empirical research.

Supervisors: José Ricardo Vaz Antunes Pereira Correia; Isabel María Parra Oller.

MSc in Corporate & Financial Risk Management, University of Sussex, United Kingdom Sep 2010 – Feb 2012

Graduate training in risk management, finance, econometrics, and economics.

BSc in Accounting & Business/Management, Brunel University of London, United Kingdom 2006 – 2010

2+2 articulation: Renmin University of China, School of International Education, Beijing (2006–2008), then Brunel (2008–2010, degree-conferring).

PUBLICATIONS & WORKING PAPERS

- **Yonghong Zhang**, Ricardo Correia, Yong Xie, and Isabel M. Parra. **CausalVerify: An Execution-Grounded Benchmark for LLM Causal Inference Workflows**. Benchmarks frontier LLMs on causal-inference workflows, pairing method/design interpretation with execution-grounded coefficient recovery. *Under review at NeurIPS 2026 (Evaluations and Datasets Track)*.
- **Yonghong Zhang**, Yong Xie, Isabel M. Parra, and Ricardo Correia. **Evidence-Grounded Auditing of Identification Assumptions in Climate-Policy Causal Evaluations**. Introduces ARGUS, a bounded LLM auditing system for DID identification assumptions using an assumption-implication-evidence rubric. *Submitted to the ACL 2026 Workshop ClimateNLP*.
- **Yonghong Zhang et al.** **The Double Measurement Confound in Agent Benchmarks: De-Scaffolding, Ground-Truth Scoring, and Reliability Beyond the Mean**. Studies how agent benchmarks can mistakenly measure scaffolds and scorers rather than model capability. *Under review at AAAI 2027*.
- **Yonghong Zhang**, Ricardo Correia, and Isabel M. Parra Oller. **Do Carbon Markets Deliver Environmental Co-Benefits? Evidence from China's ETS Pilots**. Difference-in-differences with province and year fixed effects on a 2004–2022 panel, with spatial-spillover analysis. *Working paper*.

SELECTED RESEARCH & AI PROJECTS

- **GLEAM: Multi-Agent Measurement of Firm-Level ESG Perceptions from Multilingual News**. [\[Project website\]](#) Develops a structured LLM-agent pipeline to construct firm-level environmental, social, and governance perception measures from multilingual news. Ongoing work examines their validity and association with firm-level financial outcomes.
- **ComtradeBench / OpenEnv benchmark**. Deterministic, execution-grounded environment for evaluating autonomous agents under realistic API frictions; basis for the Web Agent Track and OpenEnv Custom Track results below. demo: huggingface.co/spaces/yonghongzhang/comtrade-env

HONORS, AWARDS & COMPETITIVE ACHIEVEMENTS

- **1st Place, Best Defender, AgentBeats × Lambda Security Arena 2026**, UC Berkeley RDI × Lambda; Team MateFin; 91.4% defender win rate across 21 held-out adversarial runs.
- **1st Place (tie), Web Agent Track, AgentX-AgentBeats (Phase 1)**, UC Berkeley RDI; Google DeepMind-sponsored track; Team MateFin; prize: \$8,500 in Google Cloud credits and \$750 in Lambda credits.
- **2nd Place, OpenEnv Custom Track, AgentX-AgentBeats**, UC Berkeley RDI; Team MateFin; prize: \$3,000 in Hugging Face credits and six months of Hugging Face Pro per team member.
- **Honorable Mention (Accésit), Best Poster Competition, UAM Doctoral Week 2026**; Selected Representative Poster for the Economics and Business Doctoral Programme.
- **Third Prize, 2023 Annual Key Research Project, Rizhao Finance Society, China, awarded April 2024**. Project: Climate Change Risk and Financing Constraints: A Case Study of China.
- **Rising Star of the Year, Shandong Vocational University of Foreign Languages, 2023**.

EXPERIENCE

-
- PhD Candidate**, Universidad Autónoma de Madrid Feb 2024 – present
Doctoral research on carbon markets, climate-policy evaluation, and environmental co-benefits using panel econometrics and causal-inference designs; reproducible empirical workflows in Stata, R, Python-based evaluation workflows, LaTeX/Overleaf, and GitHub.
- Lecturer in Financial Management**, Shandong Vocational University Sep 2022 – Apr 2024
of Foreign Languages
Taught Corporate Finance and Financial Accounting (bilingual English/Chinese); mentored competition teams; served as an academic reference for students applying to postgraduate programmes in Europe.
- Founder & Owner**, Bubble-Tea & Coffee Retail Chain, China May 2014 – 2022
Built and ran an independent multi-store chain over ~8 years, including operations, finance, hiring, suppliers, and local marketing.
- English Teacher**, Beijing Road Middle School, China 2012 – 2014
Taught English to secondary-school students.

SELECTED PRESENTATIONS

-
- **Wolpertinger Conference 2026, Marco Onado PhD Session, Roma Tre University** — accepted presentation.
 - **ICFAB 2026, University of Surrey** — accepted presentation in the Sustainable Finance track.
 - **Doctoral Seminar, Universidad Autónoma de Madrid** — presented the China ETS environmental co-benefits study.

METHODS & TECHNICAL SKILLS

Econometrics: difference-in-differences, panel fixed effects, event studies, placebo tests, spatial spillovers, robustness checks. Machine learning / AI: LLM agents, retrieval-augmented generation (RAG), evaluation benchmarks, benchmark validity, evidence grounding, agentic workflows, responsible AI. Research software & tools: Python-based research and evaluation workflows, Stata, R, Git/GitHub, LaTeX/Overleaf, Hugging Face Spaces. Data domains: climate finance, carbon markets, environmental emissions, Chinese provincial panel data, multilingual news, trade/policy datasets. Communication: academic writing, research presentations, poster design, English research communication, Chinese-English translation.

TRAINING & SUMMER SCHOOLS

-
- Machine Learning Summer School (MLSS 2026), Max Planck Institute for Intelligent Systems & University of Tübingen, Germany. Selected participant.
 - Agentic AI MOOC, UC Berkeley RDI (Center for Responsible, Decentralized Intelligence): Legendary Tier (top completion tier), Fall 2025.

LANGUAGES

Chinese (native); English (professional working language); Spanish (B1, in progress).