I-han Lai

ihan.lai0924@gmail.com 650-213-7534 ihan-lai.com

Last update: Sept 2026

EDUCATION BACKGROUND

Management Science & Engineering, Stanford University Sep 2023 - Jun 2028 (Expected)
*Ph.D. major in Operation Research, Ph.D. minor in Statistics (Advisor: **Stefan Wager**)* *Palo Alto, CA*

Economics and Management, Mathematics, Tsinghua University Sep 2018 - Jun 2023
*B.Sc. in Finance and Applied Mathematics (Advisor: **Bo Li**)* *Beijing, China*

Booth School of Business, University of Chicago Sep 2021 - Aug 2022
*Research Professional (Advisor: **Dacheng Xiu, Matthew Notowidigdo, Zhiguo He**)* *Chicago, IL*

Department of Statistics, UC Berkeley Jan. 2021 - Jun. 2021
*Exchange student in Statistics (Instructor: **Noureddine El Karoui**)* *Berkeley, CA*

RESEARCH INTEREST

Causal Inference; Data-driven decision-making; Reinforcement Learning; AI Evaluation

My research develops statistical and computational methods to improve data-driven decision-making in complex, evolving systems, drawing on causal inference, decision theory, and reinforcement learning.

HONORS & AWARDS

Enlight Foundation Graduate Fellowship	2025
School of Engineering Fellowship	2024
Tsinghua Excellent Research & Academic Scholarship	2022, 2021, 2020
Baogang Scholarship	2021
National Scholarship	2020

RESEARCH

Estimating Dynamic Marginal Policy Effect Under Sequential Unconfoundedness
*with **Stefan Wager**; submitted to **Journal of the American Statistical Association (JASA)***

Abstract: We develop methods for estimating how infinitesimal policy changes affect long-term outcomes in dynamic systems. We show that dynamic marginal policy effects (MPEs) can be identified via tractable reduced-form expressions, and can be estimated under a general sequential unconfoundedness assumption. We also propose a doubly robust estimator for dynamic MPEs. Our approach does not require observing full dynamic state information (as is typically assumed for off-policy evaluation in Markov decision processes), and does not incur an exponential curse of horizon (as is typical in non-Markovian off-policy evaluation). We demonstrate practicality and robustness of our approach in a number of simulations, including one motivated by a dynamic pricing application where people use past prices to form a reference level for current prices.

Policy Learning with Strategic Agents Under Repetitive Cross-sections
*with **Yuchen Hu, Junyi Liu, Lihua Lei**; Working paper*

Abstract: We study policy learning under strategic behavior, where individuals adjust the information they report in response to an announced decision rule. Credit applicants, for example, may change their profiles to improve their chances of approval. This setting shares the feedback central to performative prediction: deploying a rule changes the data used to evaluate and improve it. Policy updates based on existing data often treat reported characteristics as fixed, overlooking how people would respond to a different rule. Repeated updates can therefore settle on a stable but suboptimal policy. We consider repeated deployments to new groups of individuals, where decisions affect both current outcomes and the conditions faced by later groups. The challenge is to learn how strategic

responses affect long-run welfare when individuals' underlying characteristics are unobserved. We propose small, temporary changes to the announced rule, randomized before individuals respond. By examining the resulting reports and outcomes over subsequent periods, we learn which policy adjustments improve welfare after accounting for the behavioral responses they induce. We show that these experiments guide repeated policy improvements while the cost of experimentation diminishes over time. In simulated screening environments, our approach captures approximately 95% of the attainable welfare gain, compared with less than 60

Active Decision Making
*with **Puheng Li, Xiaojie Mao, Renyuan Xu**; in preparation for **Operations Research***

Abstract: Experiments often inform treatment decisions, but designs that optimize global estimation accuracy need not yield treatment rules with low deployment regret. We study sequential experiments with freely observed covariates and costly outcomes, and derive a budget-constrained collection rule from a smooth surrogate for deployment regret: enrollment is proportional to the square root of decision-weighted outcome variance, with Neyman treatment allocation among enrolled units. We show that, within a broad class of adaptive designs, deployment regret is determined by treatment-effect estimation risk along the decision boundary, that the proposed design minimizes this risk, and that a feasible version with nuisance quantities estimated online converges to the optimal collection rule. Simulations illustrate the divergence between estimation accuracy and decision quality, while an application to prompt routing between GPT-4 and Mixtral-8x7B yields approximately one third of the deployment regret of uniform or variance-targeted evaluation at equal budgets.

Integrating Operational Records and Earth Observations to Reconstruct Wildfire Suppression Activities in California
*with **Qiao Kang, Avery Bick, Jessica Yu, Stefan Wager, Michael Mastrandrea, Michael Wara**; abstract submitted to **AGU2026**, in preparation for **Science Advance***

Abstract: Wildfire suppression depends on coordinated aerial and ground operations, but retrospective evaluation is constrained by fragmented data systems, inconsistent spatial and temporal resolution, and weak links among planned assignments, observed activities, and evolving fire behavior. We present an integrated framework for reconstructing suppression activities across 864 California wildfire incidents from 2017–2025, with detailed tracking of aerial and ground resources. Using large-language-model-assisted parsing, we extract aircraft and personnel assignments, resource types, operational periods, and geographic responsibilities from 3,380 Incident Action Plans. We link these planned assignments to time-stamped ADS-B aircraft trajectories, National Interagency Fire Center (NIFC) geospatial and operational records, and retardant-delivery data, then align them with terrain-corrected GOES active product and VIIRS active-fire detections, incident perimeters, weather, fuel moisture, and terrain. The resulting 15-min, 500-m data cube distinguishes planned resource assignments, observed aircraft presence, and recorded or inferred suppression actions while retaining source-level provenance and quality flags. Across the study incidents, current inventories include 20,000+ California retardant-drop records. We use the integrated record to quantify operational coverage, establishing metrics of operational efficiency and identifying where suppression effectiveness can be evaluated reliably. This framework provides a reproducible basis for post-incident learning and future counterfactual analysis of how suppression activities influence fire progression and impacts.

PROFESSIONAL EXPERIENCE

Amazon | *Applied Research Scientist Intern* Jun. 2026
Optiver | *Trader Intern* Feb. 2021

ACADEMIC SERVICE

Reviewer: *The Annals of Statistics* (2), *Biometrika* (3), *Management Science* (2).