Quinn Lanners

Seattle, WA | lannersq@gmail.com
LinkedIn | Google Scholar | Website

Summary

Research / Applied Scientist with a Ph.D. in Biostatistics from Duke University, specializing in causal inference, experimentation, and large-scale advertising ML. Currently at Upstart, driving methodology across uplift modeling, offline evaluation, online exploration, calibration, and forecasting to optimize decisions under fixed channel cadences, long-term feedback loops, and rare outcomes in dynamic environments.

Professional Experience

Upstart — Seattle, WA
Senior Research Scientist — Dec 2025 – Present
Research Scientist — Jun 2025 – Dec 2025

- Develop targeted selection logic, model overlap frameworks, and surrogate models to optimize experimentation budgets. Through strategic allocation, these frameworks achieve a 70% reduction in Minimum Detectable Effect (MDE) for rare outcomes alongside a 3% boost in campaign performance.
- Design dynamic calibration for send optimization and build exploration-based frameworks for off-policy offline evaluation, navigating the high costs of failure and strict physical constraints of the channel.
- Lead the migration of our Direct Mail uplift modeling codebase to a modern Metaflow and Python pipeline, achieving a 20x speedup in execution time and significantly reducing compute costs.

Meta — Seattle, WA
Research Scientist Intern — May 2024 – Aug 2024

- Developed a domain-adapted, model-based approach for offline counterfactual evaluation of large-scale ad-ranking models, improving pre-deployment testing of ML systems. Work accepted at the 2024 *Causality, Counterfactuals & Sequential Decision-Making Workshop* at *RecSys* and the 2024 *Conference on Digital Experimentation @ MIT*.
- Designed and ran empirical analyses comparing multiple evaluation methods, identifying conditions where the proposed approach outperforms existing benchmarks.
- Bridged observational causal inference and experimentation by enhancing offline evaluation techniques for models later tested via A/B experiments.

Duke University — Durham, NC
Ph.D. Researcher — Aug 2021 – May 2025

- Formulated a novel partial identification and data fusion framework using double machine learning to reliably estimate causal effects from both experimental and observational data under simultaneous assumption violations. Work published at NeurIPS.
- Developed an interpretable method for estimating dynamic treatment regimes that reduced the probability of an adverse medical event by over 20 percentage points when operationalized on a real-world ICU dataset. Work published at AISTATS.
- Created a variant of multi-label learning optimized for rare event settings that improved prediction of cardiovascular adverse events in pregnancy by up to 10%. Work published as a spotlight paper at UAI.
- Engineered a new method for large scale, computationally inexpensive, and interpretable causal inference that is over 100 times faster than existing benchmarks. Work published at UAI.

Optum — Eden Prairie, MN
Data Scientist — Jun 2019 – Jul 2021

- Built and deployed a multimodal time-series model assessing the risk of every development team's proposed application update.
- Detected suspicious activity as lead python programmer for fraud investigation in the payout of the U.S. Government's COVID-19 HHS CARES Act Provider Relief Fund facilitated by Optum.
- Automated development team's Jenkins environment by creating customizable `Jenkinsfiles` to deploy machine learning models to `Kubernetes`.
- Summarized bank transaction and account activity logs into feature vectors using `PySpark`, and used Spark's `MLlib` library to detect irregular account activity.

Education

Duke University — Durham, NC
Ph.D. in Biostatistics & Bioinformatics — Aug 2021 – May 2025
Dissertation: *Leveraging Machine Learning for Interpretable Causal Inference*
Advisors: Cynthia Rudin, David Page, & Alexander Volfovsky

Loyola Marymount University — Los Angeles, CA
B.S. in Applied Mathematics (Minor in Biochemistry) — Aug 2015 – May 2019
Summa Cum Laude
Thesis: *Neural Machine Translation & Deep Learning*
Advisor: Thomas Laurent

Publications

Quinn Lanners, Cynthia Rudin, Alexander Volfovsky, and Harsh Parikh. Data fusion for partial identification of causal effects. *Advances in neural information processing systems*, 38:150638–150702, 2026

Quinn Lanners*, Qin Weng*, Marie-Louise Meng, and Matthew M Engelhard. Common event tethering to improve prediction of rare clinical events. In *The 40th Conference on Uncertainty in Artificial Intelligence*, 2024

Harsh Parikh*, **Quinn Lanners***, Zade Akras, Sahar Zafar, M Brandon Westover, Cynthia Rudin, and Alexander Volfovsky. Safe and interpretable estimation of optimal treatment regimes. In *International Conference on Artificial Intelligence and Statistics*, pages 2134–2142. PMLR, 2024

Quinn Lanners, Harsh Parikh, Alexander Volfovsky, Cynthia Rudin, and David Page. Variable importance matching for causal inference. In *Uncertainty in Artificial Intelligence*, pages 1174–1184. PMLR, 2023

Mohamed A Radwan, Himaghna Bhattacharjee, **Quinn Lanners**, Jiasheng Zhang, Serkan Karakulak, Houssam Nassif, and Murat Ali Bayir. Counterfactual evaluation of ads ranking models through domain adaptation. In *Causality, Counterfactuals & Sequential Decision-Making workshop at RecSys*. arXiv:2409.19824, 2024

* Denotes co-first authorship.

Technical Skills

Languages: Python, SQL
ML & Frameworks: XGBoost, PyTorch, Scikit-learn, EconML, CausalML, PySpark, d3rlpy, Pandas, NumPy
Infrastructure & Tools: Metaflow, Databricks, Redshift, AWS, Docker, Kubernetes, Jenkins, Slurm, Git
Methodologies: Causal Inference, Uplift Modeling, Off-Policy Evaluation, Contextual Bandits, CUPED, Synthetic Controls, Instrumental Variables, Online Experimentation, Surrogate Modeling, Partial Identification