Ryan James Martin

Ryan.J.Martin1@gmail.com • (805) 698-1229 • San Francisco, CA • LinkedIn

Senior data scientist with 7+ years building production ML and causal inference for **customer lifetime value estimation, incrementality, and automated decisioning**. Expertise in forecasting, price elasticity, experimentation, and end-to-end deployment (Python/SQL, Airflow, tests/monitoring). PhD in Economics.

Relevant Skills

End-to-End Model Building: Python (scikit-learn, XGBoost, PyTorch, TensorFlow, StatsModels, dask, portkey), SQL (Snowflake/PostgreSQL), Airflow, Databricks, Docker, AWS Batch Aps, MLOps.
Causal Inference & Experimentation: A/B tests, geo-clustered XPs, switchback XPs, quasi-experiments, structural models, pair-based designs, diff-in-diff, synthetic control, propensity score matching, doubly robust estimation, meta learners, instrumental regression, regression discontinuity design.
Risk & Pricing: eligibility/scoring, risk-based pricing, elasticity, profitability tradeoffs, interpretable predictive models.
Forecasting & Reliability: time-series and distributional forecasts, uncertainty/intervals; deployment, monitoring, simulation, tests.
Search and Deep Learning: reinforcement learning, recommender systems, deep learning, transformers, embeddings, leveraging LLMs, NLP

Relevant Experience

Faire Jul 2024–Present

Senior Applied AI/ML Scientist; Brand Growth, Onbarding, Acquisition, and Success

- Serve as *data science pod lead* for the *brand growth*, *brand success*, and *brand onboarding and acquisition* teams. I have a cross-functional leadership role, among a team of engineers, a project manager, analysts, sales people, and customer support reps.
- Own, build, and iterate on a suite of ML models and airflow DAGs. These models cumulatively deliver millions of dollars to Faire every year. Key models include: (i) a lead scoring model for all inbound applicants, (ii) a recommender model that boosts exploration of underexplored brands on our new brand surfaces (iii) a "top brand" model to predict performance for the best-of-the-best brands, (iv) an "entity resolution" and lead generation model suite that leverages LLMs and retailers' spending data to generate new leads and match existing leads to additional data
- Leveraged MLOps tools and best practices to build and deploy these models in a production environment with both batched and realtime inference.
- Use econometric and ML tools to answer specialized questions like quantifying brand net-new orders vs cannibalization by category ("incrementality"), and the best brands to add to Faire.
- Design, run, and analyze experiments to guide business investments and precisely measure the value of my data science solutions.

GoPuff Dec 2023–May 2024

Senior Data Scientist; Experimentation and Pricing

- Served as *team lead* in the pricing-data-science team. Had numerous leadership responsibilities such as leading discussions on the team's roadmap formulation, providing guidance and mentorship on technical solutions, and serving as technical expert in my subject area for discussions in technical meetings with leadership, including CEO and CTO.
- Delivered multi-million dollar net revenue gains through data-driven, automated improvements to GoPuff's pricing policies. Measured improvements with A/B tests, switchback tests, diff-in-diff studies, and Matched Market testing.

- Developed machine learning, causal inference, and econometric tools in Python for GoPuff's experimentation platform and price-assignment engines as part of a dynamic team of data scientists and engineers. Pricing work included price sensitivity, subscription "fam" pricing, expiring good pricing, and delivery fee pricing.
- Ran causal inference studies to measure the impact of advertising: *Market Match Testing*
- Developed test suites to ensure the codebase is robust to continuous development.

Flexport May 2021–Dec 2023
Senior Data Scientist/Economist

- Served as the team lead of Pricing and Causal Inference data scientist working group. Coordinated workstreams within my team and across the other data-science teams to solve key business problems and meet high-impact business goals.
- *Used machine learning, causal inference, and econometric tools* to build end-to-end solutions in Python and SQL such as (i) customer segmentation; (ii) predicting customer's lifetime values and churn probabilities; (iii) estimating heterogeneous price elasticities; and (iv) estimating the value of faster shipping
- Performed causal inference studies to help guide business decisions and investment, working closely with CFO's team. This included explaining what parts of our platform made each of our customer segments the stickiest, how to optimize sales incentives, and how to price novel offerings like express ocean freight.
- Measured impact through A/B tests and geo-clustered experiments.

Bank of Canada Jul 2019–May 2021
Research Scientist (Senior Economist)

- Built machine-learning and time-series models to automate bank-note-demand (-volume) forecasts and report bank-note-demand volatility (risk).
- Built and estimated sophisticated econometric and statistical models in Python to measure Canadian demand elasticities (or price sensitivities) to payment methods at the point of sale.
- Continued graduate research on estimating search ranking's impact on customer demand and willingness to pay.

Microsoft Jun 2016–Sep 2016
Research Intern, MSR

- Queried massive datasets on Azure with SQL, python, and C#.
- Applied knowledge of web scraping, string parsing, and Natural Language Processing (NLP) to build and estimate data features for inputs into machine learning and econometric models in python.
- Used econometric models to gain deep insights into high-impact research questions around customers' online search behavior and purchase decisions.

Education

University of California, Los Angeles (UCLA)
PhD, Economics Aug 2013–Jun 2019
Focus: *Causal Inference, Industrial Organization, Marketing, ML and Econometrics*
Research Topics: *Demand estimation with costly search in online marketplaces, Demand estimation on healthcare exchanges, costly search and pricing competition, optimal reserve prices in timber auctions*
Additional Graduate Training: Statistics (9 Courses) Spring 2016–Fall 2017

University of California, San Diego (UCSD)
BS, Probability and Statistics, *magna cum laude* Fall 2010–Summer 2012