Joseph Wang · 汪久量

Data & Business Analytics | Customer Analytics · Forecasting · Pricing & Inventory

+86 137 5718 7261 · joseph@josephjwang.com · LinkedIn · GitHub · Portfolio

Shanghai, China · English (fluent) · Mandarin (native)

SUMMARY

Analytics professional with 10 years of experience, six in North America (GE, Mercedes-Benz Canada, Lordco). Pilot categories saw gross margin rise 2–4 pp versus pre-change levels under my pricing recommendations; a three-arm randomized test raised redemption 6.2 pp; a repurchase model reached 0.78 AUC out of time. Use SQL, Python, Tableau, and Power BI to build metrics, experiments, and forecasts for business decisions. Fluent in English and Mandarin.

CORE SKILLS

Data processing SQL · Python (pandas, scikit-learn, statsmodels) · data cleaning and reconciliation

Business analytics Metrics · funnels and cohorts · P&L · pricing and inventory · Tableau · Power BI

ML & forecasting Regression · XGBoost / LightGBM / CatBoost · demand forecasting · CLV · temporal validation

Experimentation A/B and multi-arm tests · sample size / MDE · confidence intervals · guardrails · holdouts

Applied delivery Streamlit · Git · governed ChatBI · LLM evaluation · QLoRA · bilingual reporting

PROFESSIONAL EXPERIENCE

Yuyue Information Technology · Business Analytics & Data Science Consultant — Shanghai / Canada · Sep 2022 – Present

Provide analytics consulting for an employee-benefits platform and SMBs, covering requirements, data preparation, modeling, and delivery. The main project spans about 120K voucher records over three years, 20K end users, and 80+ entitlement SKUs. Delivered three other SMB projects on dashboards, sales and collections, and inventory aging.

- **Randomized experiment.** Designed and ran a three-arm test among about 1,500 users who had visited the platform but had not redeemed their benefits (~500 per arm). Fourteen-day redemption was 34.8% with reminders versus 28.6% in control (+6.2 pp). The test detected no additional lift from rewards; recommended prioritizing low-cost reminders.
- **Metric framework.** Integrated data across six voucher lifecycle stages with explicit table grains and reconciled definitions. Built a 32-metric dictionary and consolidated eight manual reports into one dashboard for operations, procurement, and leadership reviews.
- **Forecasting & resource planning.** Built a four-week rolling forecast by week and entitlement category: out-of-time WMAPE of about 22% versus 28% for a moving-average baseline. Translated forecasts into procurement, budget burn, and expiry exposure; compared allocations using auditable rules and scenarios under budget, inventory, and fulfillment constraints.
- **Governed ChatBI.** Built a prototype mapping about 20 recurring questions to reviewed SQL templates, with parameter and result checks, permission filters, and out-of-scope refusal. On a standard question set, turnaround fell from 1–2 business days of manual pulls to minutes.

Lordco Auto Parts · Pricing & Inventory Analyst — British Columbia, Canada · Jul 2020 – Jun 2022

Western Canadian auto-parts retailer and distributor. Responsible for pricing and inventory analytics across 100K+ SKUs, multiple stores, and Retail / Trade / Fleet customer channels.

- **Pricing recommendations & pilots.** Recommended prices for category-manager approval under cost, target-margin, competitor-gap, and price-change constraints. Pilot categories recorded a 2–4 pp increase in gross margin versus pre-change levels.
- **Operating metrics.** Joined POS, product, cost, competitor, inventory, and supplier data into weekly SKU × store × customer-channel monitoring, covering 90%+ of active SKUs and tracking margin, units, weeks of supply, and aging risk.
- **Demand modeling.** Used price, inventory, seasonal, and product features to benchmark linear models against XGBoost / LightGBM / CatBoost by category, and simulate volume and margin at candidate prices.
- **Price–inventory rules.** Flagged slow-moving, high-stock products for promotion or clearance and high-demand, low-stock products for replenishment and discount limits. Automated weekly reports, saving about one analyst-day per week.

Mercedes-Benz Canada · Business & Retail Planning Analyst — Toronto, Canada · Nov 2017 – Jun 2020

Customer lifecycle and dealer performance analytics for the national retail network, spanning new-vehicle retail, finance and lease, and after-sales.

- **Customer lifecycle.** Combined CRM, vehicle, purchase and finance contracts, repair orders, and satisfaction data into purchase-quarter cohorts. Tracked 6/12/18-month service return and contract-expiry windows to identify attrition, certified pre-owned, and renewal or repurchase opportunities.
- **Propensity model.** Built a 12-month repurchase and lease-renewal model with logistic regression and gradient boosting, reaching about 0.78 AUC out of time. Combined scores with contract-expiry windows to publish monthly ranked customer lists piloted with 12 dealers.
- **Dealer performance.** Translated sales, after-sales, and customer-experience targets into 30+ KPIs. Built Tableau reporting for incentive calculation, monthly tracking, and rebate settlement across 50+ dealers, saving the team about two days a month.

Earlier experience: GE Appliances, Financial Analyst, Ontario, Canada (Jul 2016 – Nov 2017). Category budgets and monthly rolling forecasts; price-volume-mix analysis of margin variance to support pricing and promotion decisions. MetLife, Actuarial Analyst Intern, Shanghai (Jan 2013 – Apr 2013). Life-insurance pricing and reserve analysis; SQL data extraction and validation, Excel actuarial models, and VBA automation.

SELECTED PROJECTS

Personal projects on public data; live demos at josephjwang.com/demos

Customer Intelligence Platform | RFM · CLV · churn · CRM actions

- Built RFM segments and BG/NBD + Gamma-Gamma models on 805K retail transactions from 5,878 customers, using an 18-month calibration period and six-month holdout. Customer-level predicted and actual holdout revenue had a correlation of 0.84.
- Diagnosed poor churn ranking from 1−P(alive), which scored one-time buyers as low risk (AUC 0.44). Replaced it with supervised logistic regression reaching 0.81 AUC out of time, and mapped customer segments to retention and win-back recommendations.

Review Intelligence | LLM evaluation, fine-tuning & deployment choice

- Compared lexicon baselines, zero-shot and few-shot LLMs, and QLoRA fine-tuning on 46,730 reviews across 18 aspect categories, with an external English-data check. Best Chinese sentiment Macro-F1: 0.757. Evaluated deployment options on structured-output validity, latency, and cost.

EDUCATION

Northwestern University · M.S. Data Science · GPA 4.0 / 4.0 — USA · Jun 2022

- **Industry capstone.** Built a Ridge guideline-pricing model on 911K B2B telecom quotes, generating recommended prices for 336K open opportunities. Compared gradient boosting and selected pricing logic that sales teams could audit; delivered win/loss and competitor-trend dashboards.

University of Waterloo · B.Math, Double Major in Actuarial Science & Statistics — Canada · Jun 2015