Annapragada Vamsy Vrishank

mail@vamsy | +1 (201) 936-2248 | New York, NY | vamsyvrishank | vamsyvrishank | vamsy.org

Education

Stevens Institute of Technology, Hoboken — New Jersey, United States
Masters in Financial Engineering (Specialization: Algorithmic Trading); CGPA: ***3.959/4.0*** — *Jan 2025 – May 2026*

National Institute of Technology (NIT), Trichy — Tamil Nadu, India
Bachelors in Metallurgy, Materials Science and Microeconomics; CGPA: ***8.1/10.0*** — *July 2017 – May 2021*

Experience

PayPal — Bangalore, India

Software Engineer II (Customer Data Platform) — *July 2024 – Jan 2025*

- Upgraded batch Airflow ETL pipelines into NRT event-driven streaming (Java, GCP Pub/Sub, Dataflow) across **500K+ daily transactions**; cut data latency from **6 hrs to 10 min** for fraud monitoring and analytics teams.
- Unified **200M+ customer profiles** across Braintree, Venmo, and Hyperwallet; normalizing raw events into common schema using probabilistic identity resolution; golden records loaded into BigQuery served risk, fraud and ML downstream

Software Engineer I (Enterprise Data Infrastructure) — *Sep 2021 – July 2024*

- Re-platformed **800TB+** enterprise infrastructure from on-prem Hadoop/Spark to GCP; managed **200+** production ETL/ELT pipelines across Spark and Pub/Sub.
- Migrated Braintree's data warehouse on BigQuery with query-pattern-aligned partitioning; reporting at **60% lower cost**, saving **$4M annually**; designed fact and dimension tables for analytics and research consumers.
- Designed ETLJobsServ to solve connection exhaustion across 300+ concurrent Airflow pipelines by replacing ephemeral per-pipeline connections with a long-running Spring Boot service

Projects and Research

- **Execution-Aware Alpha Modeling (Stevens Capstone)**: Joint optimization of decaying alpha signals and execution schedules using OU process augmented Almgren-Chriss framework; derived closed-form optimal trading schedule governed by dimensionless coupling ratio $\Phi = \theta/\kappa$. Statistically validated that joint optimization outperforms standard AC in the coupling regime ($0.5 \leq \Phi \leq 2$) via paired t-tests.
- **StoryForge: Multimodal Agent**:Agentic pipeline (Google ADK, Gemini 2.0 Flash) chaining 3 LLM agents; RAG-style context injection and prompt engineering across agent handoffs; statistically grounded evaluation framework measuring output coherence and model quality. FastAPI SSE backend; Dockerized on GCP Cloud Run.
- **Adaptive Volatility Regime-Based Execution Framework**: 3-state HMM with HAR-RV-J realized variance forecasting decomposing equity dynamics into diffusion and jump components; integrated Lee-Mykland jump detection and Hawkes flash-crash modeling for regime-conditioned signal generation. CVaR-constrained VWAP/TWAP execution sizing calibrated to Almgren-Chriss; Sortino 2.41 vs Sharpe 1.57; walk-forward validated on out-of-sample data.
- **Stablecoin Microstructure and Liquidity Fragmentation (IAQF 2026)**: 1.8M+ tick-level pipeline across 3 exchanges; implemented Kyle's λ, Amihud illiquidity, Roll spread, and HHI on noisy cross-exchange datasets. SVB crisis as natural experiment; Roll spreads 20x, price dislocation 15x; identified quote-currency layer as primary shock transmission channel.
- **Equity Derivatives Pricing Engine**: Implemented pricing models in C++ with Monte Carlo simulation,PDE solvers; computed implied volatility surface via Newton-Raphson inversion Factory/Strategy architecture.

Skills

Languages & Libraries: C++, Python (Pandas, NumPy, FastAPI, ML Libraries), Java (Spring Boot), SQL
AI/ML & Quantitative Methods: Machine Learning, Time Series Econometrics, Hidden Markov Models, Stochastic Calculus, Options Pricing, Convex Optimization, Backtesting
Data & Distributed Systems: Data Pipeline Architectures, Apache Kafka (Messaging Systems), Apache Spark, GCP, AWS, PostgreSQL (Relational DBs), Redis (Non-relational DBs)
Systems & Architecture: Unix/Linux, Windows, Docker, Kubernetes, Microservices, Design Patterns, REST APIs, Distributed Systems