Emergent IV
Research & Development Broad Agency Announcement

PETABYTE ECONOMICS CORP.

DLA EMERGENT IV R&D BAA

White Paper Cover Sheet

Project Title	Decision Theory Machine Learning for Demand Forecasting & Inventory Optimization
Proposer Organization	Petabyte Economics Corp. UEI: VRNNE19YME58 (U.S. Small Business)
BAA Area of Interest (from Attachment B):	H: Artificial Intelligence/Machine Learning Applications 2: Demand forecasting & inventory optimization
Technical Point of Contact (POC)	[redacted]
Administrative POC	[redacted]
Other Team Members (subcontractors and consultants), if known/applicable	[redacted]
Estimated Total Cost (Base + Options)	[redacted]
Estimated Period of Performance	[redacted]
Identify any other solicitation(s) to which this concept has been proposed	This concept is related to NSF SBIR project pitch 00077364.

Executive Summary

Decision Theory Machine Learning for Demand Forecasting & Inventory Optimization

There is a right way for the Defense Logistics Agency (DLA) to do inventory placement. Over 70 years ago, a Nobel Prize-winning economist and his coauthors proved that if the demand distribution and the inventory costs and benefits are known, there is an optimal solution to the so-called "newsvendor problem" (see Arrow et al., 1951). Until recently, however, this solution wasn't tractable at the scale and complexity needed for modern supply chains.

We've changed that. Our decision-theoretic machine learning method, based on [redacted], can compute the optimal inventory placement for millions of line items and thousands of storage and distribution centers to support demand coming day by day for years.

Our inventory solution builds on a Bayesian forecasting system that can incorporate trend breaks, expert opinions, and external forecasts, including those created using generative AI, to produce well-calibrated, highly accurate forecast ensembles that properly reflect future uncertainty. It uses an interpretable, white-box economic model to empower DLA personnel to understand their demand's likelihoods and conduct counterfactual analyses. Moreover, its forecast ensembles are consistently aggregable, so they can be paired with loss function information that captures DLA's posture and precision to minimize its supply chain risks.

Petabyte Economics Corp. (a US small business) is a private R&D lab in Vermont. It is led by Dr. Ed Egan, a computational economist specializing in the R&D and commercialization of big data machine learning systems. [redacted]

We are currently working on two products: "CatFish" (see Egan, 2023a & b), our Bayesian demand forecasting system that is in alpha testing, and "Newsvendor" (see Egan, 2024), our decision-theoretic method that is in R&D. We successfully pitched the National Science Foundation's SBIR program earlier this year (see NSF, 2024) to supplement CatFish's functionality and develop the machine learning methods behind Newsvendor.

We are eager to partner with DLA to deploy CatFish on its Azure cloud, build necessary data pipelines, and experiment to customize its model and enhance its functionality to meet America's combat support needs. We would work with DLA to collect and incorporate your experts' decisions and beliefs and deduce appropriate loss function data from DLA supply chain management systems. Then, we would conduct the R&D necessary to create and deploy an operational version of Newsvendor, refining outputs from CatFish and combining them with loss function information to produce posture and precision appropriate optimal inventory assignments.

Emergent IV
Research & Development Broad Agency Announcement

Project Summary

Project Title	Decision Theory Machine Learning for Demand Forecasting & Inventory Optimization
Proposer Organization	Petabyte Economics Corp. UEI: VRNNE19YME58 (U.S. Small Business)
BAA Area of Interest (from Attachment B):	H: Artificial Intelligence/Machine Learning Applications 2: Demand forecasting & inventory optimization
Technical Point of Contact (POC)	[redacted]
Administrative POC	[redacted]
Other Team Members (subcontractors and consultants), if known/applicable	[redacted]
Estimated Total Cost (Base + Options)	[redacted]
Estimated Period of Performance	[redacted]
Identify any other solicitation(s) to which this concept has been proposed	This concept is related to NSF SBIR project pitch 00077364.

Goals and Impact

Decision Theory Machine Learning for Demand Forecasting & Inventory Optimization

Our goal is for the Defense Logistics Agency (DLA) to make scientifically sound, data-driven, and human-centric demand forecasting and optimal inventory placements using modern decision theory. Our research has made this possible at scale. Specifically, we propose that the DLA create demand forecasts using Bayesian machine learning (ML) to infer the full joint distribution of demand across multiple dimensions, including items, locations, and time, and apply loss functions to these demand forecasts to optimize inventory accounting for uncertainty, costs and benefits.

The main quantitative impacts of this approach are improved accuracy relative to state-of-the-art neural net (NN) based systems (we can achieve the combined accuracy of the best NN methods enhanced by outside information), full decision-theoretic handling of uncertainty, and posture (i.e., position on the continuum of conflict) and precision (i.e., readiness-cost balanced) appropriate optimal placement. The DLA's supply chains will then provide greater national security through higher rates of first-time placement success, faster fulfillment, and better user experiences. These improvements will lower shipment and storage costs, save taxpayer money, and reduce pollution while increasing resiliency and agility.

Although the quantitative impacts are substantial, the qualitative impacts of using Bayesian decision theory are truly transformative. This approach provides critical functionality that black-box NN-based demand forecasting systems cannot (see Rudin, 2022). Specifically, it empowers human operators to create and adapt forecast ensembles that incorporate expert opinions and (even previously unseen) trend breaks, understand their demand's determinants and forecast outcome likelihoods, and conduct on-demand, mission-specific counterfactual analyses.

Our work can also be applied to supply. The US would like to forecast and assess risks in critical mineral supply (see US President, 2021). In May 2024, we discussed this with K. Watt Lough, Director of Strategic Planning and Market Research for DLA Strategic Materials.

Technical Plan

The proposed work is to create an at-scale application of decision theory to enhance DLA's real-world supply chain capabilities. Specifically, we would research, develop, and deploy a(n):

1. **Forecasting system**: A human-centric, aggregably consistent, Bayesian multidimensional, interpretable, white-box economic model of demand that allows coherent operational, strategic, and financial planning within and across levels of analysis from the same forecast model.
2. **Action-contingent model**: Methods to translate multidimensional forecast ensembles into a model integrating expert opinion, trend breaks, regime changes, and other decision inputs at appropriate grains (See [redacted].
3. **Optimal inventory placement system**: Integrate loss function information into the action-contingent model and optimize inventory accounting for uncertainty, precision and posture/scenario using a decision theoretic method based on [redacted]

The demand forecasting system will leverage any time-series model (for examples, see Godahewa et al. 2021), including generative AI models, at any grain, to capture complex dynamics and ensure upgradeable best-in-class accuracy without compromising on human-interpretable intelligence. It will also use locational and other covariates to increase demand forecasts' accuracy and calibration and allow DLA to produce scenario analyses.

The optimal inventory placement system will use decision theory to make transparent, data-driven, human-centric recommendations: Different items have different (potentially time varying) forecast levels and uncertainty, and different (potentially scenario varying) asymmetric costs and benefits in different locations and postures, so need different optimal inventory placements to minimize risk. Human operatives can review and interactive with this logic and its underlying data.

Current state-of-the-art commercial supply chain systems often use NN-based artificial intelligence (AI) to create item-time placement recommendations. NN-based AI approaches generally scale, have few context-specific requirements, and need little end-user expertise. However, because they are practical black boxes with no theoretical grounding, they defy human interpretation, which fatally compromises their use in effective decision making. NN models also do not produce predictive distributions and their forecasts are not aggregably consistent.

The Bayesian decision-theoretic approach does not have any of these limitations. Its primary challenges instead concern scalability, generalizability (each model requires considerable bespoke implementation and functionality), and manageability. We have solved the scalability issue with proprietary parallel cloud-computing-based ML methods. Addressing the generalizability and manageability issues requires R&D but also creates implementation-specific economic understanding that is crucial to adaptation and resilience.

We envision the following milestones and timeline:

- **3 months**: Report on data availability and project scope, including dimensions for the model (i.e., items, services, locations, and times), specification of inputs and outputs, major system architecture choices, etc. Review rival model(s) and determination of forecasting performance metrics (e.g., quantile loss).
- **8 months**: Deployment of forecasting system with basic functionality. Goal: competitive performance with rival system/method.
- **10 months**: Inclusion of time-trends, extended functionality, expert beliefs, and action-contingent model. Goal: outperform rival system/method at decision grains.
- **18 months**: Deployment and back-test of inventory system using historic placement and DLA loss function data. Goal: posture and precision optimized inventory placement.
- **24 months**: DLA personnel use the system to generate de novo forecasts and posture-appropriate inventory placement schedules.

Capabilities and Management Plan

Petabyte Economics Corp. ("Petabyte") is a US small business specializing in the R&D of machine learning methods to support cloud-scalable economic models. The key personnel are

[redacted]

Cost and Schedule

The cost and schedule depend heavily on the size and complexity of the forecasting and optimal placement system and the availability and readiness of data and computing resources at DLA needed for development. This proposal assumes R&D of a moderate scale and complexity system with comparatively little non-core development and essentially all required inputs and infrastructure, including pre-existing lakes of suitable data, will be provided by DLA.

The schedule is provided in the milestones above. The cost estimate assumes a January 1st, 2025 start, so year one spans fiscal years 2025 and 2026, and year two spans 2026 and 2027.

[redacted]

Bibliography

Arrow, Kenneth J, Harris, T., and Marshak, Jacob (1951), Optimal Inventory Policy, *Econometrica*. https://www.jstor.org/stable/1906813

Egan, Edward J. (2023a), The CatFish Handbook, Petabyte Economics Corp. www.petabyteeconomics.com/files/CatFish-Handbook.pdf

Egan, Edward J. (2023b), CatFish: Meaningful Demand Forecasts at Any Grain, Petabyte Economics Corp. https://www.petabyteeconomics.com/catfish.html

Egan, Edward J. (2024), Seeking Letters of Support for Newsvendor's NSF Grant Application, https://www.petabyteeconomics.com/files/NewsvendorPRFAQ_20240802.pdf

NSF (2024), A Supply Chain Decision Support System, National Science Foundation Small Business Investment Research Project Pitch 00077364, submitted to the Advanced Systems for Scalable Analytics (AA) on 2/26/2024. https://www.petabyteeconomics.com/files/Submitted SBIR Pitch - Supply Chain Decision Support System.pdf

[redacted]

Godahewa, R., Bergmeir, C., Webb, G. I., & Hyndman, R. J. (2021). Monash Time Series Forecasting Archive. ArXiv. https://arxiv.org/abs/2105.06643

Rudin, Cynthia (2019). Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature machine intelligence*, 1(5), 206–215. https://doi.org/10.1038/s42256-019-0048-x

U.S. President, Executive Order, "America's Supply Chains, Executive Order 14017 of February 24, 2021," Federal Register 86: 11849, document 2021-04280. https://www.federalregister.gov/documents/2021/03/01/2021-04280/americas-supply-chains

Curriculum Vitae Links:

[redacted]