



## Agentic AI for Demand, Credit, and Inventory Intelligence in Food Supply Chains

**Ghatoth mishra**

Independent Researcher

ghatothmishra@gmail.com

### Abstract

A multi-cloud big-data architecture delivers predictive services for demand forecasting, credit risk modeling, and supply chain inventory financing. Strategic partnerships provide a comprehensive range of data for a national food service domain, supported by data governance and management protocols. Predictive services facilitate collaborative demand reconciliation across heterogeneous cloud environments, deepening data interaction among competing firms, while optimizing Steering Committee members' inventory levels and exposure to credit risk from suppliers, distributors, and retailers. Agentic Artificial Intelligence, integrating the principles of Autonomous Analytics and a multi-agency governance framework, performs the groundwork for decision signals regarding all three services.

While agentic artificial intelligence (AI) enables autonomous data preparation, predictive analytics, and deep learning, it does not remove humans from the loop. Instead, agentic AI facilitates deeper human-AI collaboration across the multi-cloud big-data architecture, with the agentic AI agents charged with examining results and providing interpretable explanations. Although solutions are focused on a national food service supply chain formed by the collaboration of food manufacturers, distributors, and both commercial and governmental businesses, the framework is readily applicable to other large supply chains requiring data preparation, predictive demand signals for collaborative reconciliation, risk assessment, and inventory financing facilitation.



**Keywords :** Predictive Demand, Credit Risk, Inventory Financing, Supply Chain Data Ecosystems, Agentic AI, Autonomous Analytics, Multi-Cloud Big Data, National Food Service Supply Chains, Agentic AI-Driven Multi-Cloud Big Data Architecture for Predictive Demand, Credit Risk, and Inventory Financing in National Food Service Supply Chains.

### 1. Introduction

Food service supply chains suffer from opaque, asymmetric information flows that lead to high levels of waste, stockouts, and credit risk. Data-driven architectures provide a solution through data sharing, integration, and predictive analytics. Autonomous, multi-cloud big data architectural models reducing capital expenditures make it increasingly feasible for small and medium enterprises in food-service supply chains to join a shared data ecosystem. Leveraging predictive demand and risk models, lenders and inventory holders can increase revenues by financing inventory replenishment and turnover. Yet enhancing automatic recommendations requires combining data-driven and

domain-specific knowledge. This article describes a multi-cloud big data architecture enabling agentic AI for predictive demand forecasting, credit risk modeling, and advanced inventory financing.

System components are designed and developed separately, with an emphasis on the supporting multi-cloud architecture for AI-based predictive demand forecasting and credit risk modeling. Agent roles, coordination protocols, and control mechanisms are specified for operationalization in agentic AI, enabling automation of analytics as well as operational decision-making, compliance monitoring, and data privacy management.



**Fig 1: Inventory Financing in National Food Service Supply Chains**

### 1.1. Background and Significance

The food service supply chain delivers prepared food and drink to consumers. Demand flows into this supply chain from households, businesses, health care facilities, schools, prisons, and others. Demand prediction has become increasingly important for food service distributors and their suppliers. National demand forecasts typically span much of the United States food service supply chain from consumers to manufacturers but are not yet sufficiently granular for day-to-day decisions. Predictive models must operate and be updated regularly to capture seasonal and special event variations. Within this context, a novel AI (Huang et al., 2022)–driven, agentic, multi-cloud architecture for predictive demand – estimating future deliveries to food service distributors operating in the United States – combined with predictive models for credit risk and inventory reconciliation financing signals are merged into a cohesive operational framework.

Data from 76 supplier, distributor, and retailer locations were harnessed for the credit risk scoring models deployed at food service distributors. The accuracy of these scoring models, along with ongoing monitoring processes to capture any material changes in the risk factors, will ensure reliability when delivering credit risk financing signals to distributors' suppliers. Collectively, this novel data-driven architecture contributes to the scientific understanding of

predictive demand in food service supply chains, enhancing network resilience through collaborative demand and supply management.

| Model       | RMSE   | MAE    | MAPE   | % MBE  |
|-------------|--------|--------|--------|--------|
| Naive (t-1) | 13.213 | 11.264 | 12.043 | -0.359 |
| MA(4)       | 8.786  | 6.755  | 7.430  | -0.492 |
| Better      | 10.409 | 8.812  | 9.697  | -0.598 |

**Table: Metrics table (synthetic example)**

### 1.2. Research design

An experimental proof-of-concept architecture deploys novel agentic AI, predictive demand analytics, and credit risk modelling on a secure multi-cloud big data platform integrated with the national banking sector, enabling predictive financing signals for inventory stock replenishment in food service supply chains. A heterogeneous data ecosystem incorporates structured operational data and unstructured social media narratives from the supply chain, food and beverage, and credit risk domains. Predictive demand models forecast next-week consumption at food service restaurants at the national city level. Following inventory reconciliation, credit provisioning scores are derived for all borrowers at the individual outlet of a national food service supply chain, enhancing bank–client relationships during COVID-19 lockdown conditions. Potential deployment in an agentic AI-controlled Secure Multi-Cloud Multi-Tenant Multi-Organizational Data Lake Platform is also discussed.

Multi-cloud big data infrastructures combine varied services from different providers to create complex ecosystems for data integration, governance, and operationalization. Regionally located on-premises preparation security guards establish a new set of data operations and agents at the boundaries of the system, enabling data to enter and leave



data centers without modification while keeping cloud-storage loading de-coupled. Mission-oriented consensual agentic AI governs ecosystem data traffic and orchestrates the integrated feeding of heterogeneous AI-ML models deployed on PaaS clouds. Autonomous ML enables cloud services to function without human intervention.

## 2. Theoretical Foundations and Context

Multi-cloud big data architectures make it possible to cluster synergistic analytical workloads while exploiting the capacity of multiple cloud service providers (CSPs). Unlike hybrid solutions that strive for a single machine—on premise or in the cloud—Agentic AI enables a new way to orchestrate end-to-end analytics as a data ecosystem of consumer and producer agents embedded within autonomous-systems principles. Such architecture introduces a novel AI layer above the traditional infrastructures, services, and datasets in a dedicated analytics cloud powered by agentic technology.

Empowered with autonomy, agency, intelligence, and sustainability, the analytics agents perform data-centric tasks with a high degree of business and technical independence. Data science workflows are automatically initialized from the scheduled ingestion, monitoring, and consolidation of the required datasets. Within the dedicated cybersecurity perimeter, agents and data products comply with the legislation and data policies on privacy and security.

### Equation 1: MBE (Mean Bias Error)

Step-by-step:

1. Compute errors  $e_t = y_t - \hat{y}_t$
2. Average:

$$MBE = \frac{1}{T} \sum_{t=1}^T (y_t - \hat{y}_t)$$

Interpretation:

- $MBE > 0$ : predictions tend to be **too low** (under-forecast)
- $MBE < 0$ : predictions tend to be **too high** (over-forecast)

**Goal:** penalize large errors more strongly

Step-by-step:

1. Errors  $e_t = y_t - \hat{y}_t$
2. Square errors  $e_t^2$
3. Average  $\rightarrow$  MSE:

$$MSE = \frac{1}{T} \sum_{t=1}^T (y_t - \hat{y}_t)^2$$

4. Take square root  $\rightarrow$  RMSE:

$$RMSE = \sqrt{MSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T (y_t - \hat{y}_t)^2}$$

**Goal:** error as a percentage of actual

Step-by-step:

1. Absolute error  $|y_t - \hat{y}_t|$
2. Divide by  $|y_t|$  (guard against 0)
3. Average and convert to percent:

$$MAPE = \frac{100}{T} \sum_{t=1}^T \left| \frac{y_t - \hat{y}_t}{y_t} \right|$$

### 2.1. Multi-Cloud Big Data Architectures

Recent years have seen the emergence of an increasing





number of commercial and open-source platforms for harnessing big data, powered by multifarious, cloud-based technologies. Nevertheless, organizations may be more concerned about the functionality or operational capabilities of such platforms than the actual architecture. Multi-cloud big data architectures utilize more than a single cloud platform and comprise three distinct structure layers: data management, project management, and application management. Recent developments in Big Data Science indicate that cloud singularity is no longer relevant when it comes to data storage, processing, and analytics requirements. Phases of projects require heterogeneous environments, including the use of multiple object stores, RDDs, streaming services, and orchestration services, across a range of vendors.

Architectures must ensure that all needed management environments are created in agreement with the respective Big Data Science planning. Any failure or omission can lead to unnecessary overhead delays for projects. In an active, dynamic environment for data ingestion, cloud redundancy is crucial for greater reliability of the data for further operations. A multi-cloud architecture guarantees availability, although the cost must be closely monitored. Agent-based innovation and digital twin offerings present training agencies with a fertile ecosystem in which to develop proprietary training data for their services. The Multi-cloud layer orchestrates Cloud Unity to play the pivotal role as Captain of the Cloud for each Cloud Service Unit and Data Distribution Unit, defining and managing all services in the data consumption and data creation phases of the Data Project Life Cycle.

## 2.2. Agentic AI and Autonomous Analytics

Increasingly, autonomous technologies driven by intelligent agents represent a new frontier of artificial intelligence (AI) applications. Agentic AI encompasses the virtual agents that govern web activities of many organizations (Buckley and Shankar, 2022), such as negotiating and transacting in electronic marketplaces and coordinating robots in supply chain and logistics operations. Autonomous systems employ

decision-making technologies with zero or little human intervention, including self-driving cars, drone swarms, and automated missile defence. Underpinning both types of agentic AI are intelligent decision engines that continuously gather data, identify patterns, recognize changes in environments, and suggest courses of action. Autonomous analytics sustains such operations through the automated conduction of analytical tasks over large datasets. Although AI has long enabled autonomous and semi-autonomous analytics, both deeply embedded into data-intensive application domains like cybersecurity and financial fraud detection, a new class of agent-dominated architectures is imposing fresh requirements on these services.

The development of analytical capabilities is thus being performed under the demand of agent systems and services that independently collaborate with other agents as well as with humans to perform specific coordination and control tasks. Similar to autonomous analytics, agent-enabled analytics is focused on the automated performance of large-scale analytical activities in the background, imposing limited or controlled requirements on developers yet addressing demand arising from agent infrastructures.

## 3. System Architecture and Data Ecosystem

Agentic AI–Driven Multi-Cloud Big Data Architecture for Predictive Demand, Credit Risk, and Inventory Financing in National Food Service Supply Chains

A national multi-cloud big data architecture enables predictive analytics for demand, credit risk, and inventory financing by food services. Opportunities and challenges arise from a greater reliance on multichain ecosystems and the need for interoperability among providers, distributors, and customers of supply chain ecosystems. A prototype multi-cloud architecture consisting of public platforms and private environments supports the twenty-four-seven integration of data originating from quasi-independent



ecosystems dispersed nationwide. Support from the associated data governance framework mitigates privacy and security risks in data deployment.

Big data analytics requiring artificial intelligence, machine learning, and data mining have entered the operational stage. Yet their deployment remains limited and challenged by vertical instead of horizontal solutions confined to single cloud environments. Agentic AI-driven autonomous analytics and control seek to address these limitations by deploying business models that continuously monitor identified patterns and signals to facilitate operational decision-making and execution with minimal human intervention. Essentially, the analytics procedure is repeatedly undertaken by agents in accordance with predefined rules and protocols. Heterogeneous cloud environments are connected by solutions such as Apache NiFi that continuously ingest data, while agentic AI brokers supervise the orchestration and monitoring of the multiservice multichain implementation.




Fig 2: System Architecture and Data Ecosystem

### 3.1. Data Sources and Data Governance

The described architecture exploits additional data from diverse external sources to enhance decision-making. Food safety officials provide predictive hazard data for restaurants, hotels, and central kitchens in municipalities.

Banks contribute encrypted risk data for potential inventory-financing transactions. Credit-reporting bureaus disclose proprietary default-risk data from the restaurant and hotel sectors. Furthermore, a leading provider specializes in financial-and-marketing analytics, disclosing business-data packets for supply-chain risk and promotion analyses.

Public policy and societal goals, such as equity and sustainability, influence the governance of the additional data. Appropriate regulatory frameworks discuss data accessibility, risk share, use permissions, supervision, and integration to ensure that the value justifies the privacy-and-security restrictions imposed on third parties. Trust frameworks support these initiatives while considering the apparent and potential data deliveries. Restrictions imposed by GDPR, CCPA, PIPEDA, and other privacy-related laws determine the data sources admissible to each data group, and custom enabling trusts ease the friction of trustlessness across the cloud.



### 3.2. Data Ingestion, Processing, and Storage

Five categories of analytics services are required: data ingestion; data fusion, preparation, cleaning, and transformation; data storage and indexing; data provenance; and data documentation. Following the data governance



framework, the data ingestion agents collect data in accordance with the respective service-level agreements. Collected data are temporarily stored in staging areas before undergoing automated quality checks. Whereas failure of these checks will cause data to be quarantined until resolved, passing data will be redirected to appropriate onboarding processes.

Data-onboarding agents orchestrate the subsequent treatment of incoming data from the staging areas. This includes applying transformations to yield analytic-ready datasets; audio transcriptions; data assemblage, fusion, and linking; and indexing for rapid querying by analytics engines. Onboarding agents also handle data provenance recording and metadata population for documentation and indexing agents. Core data assets are then persistently stored in specific cloud storage zones for long-term retention. Each zone has a different purpose and quality threshold: raw zones hold all categorical data, including poor-quality batches; refined zones keep data sets that have passed quality-assurance protocols; and gold zones retain data assets deemed relevant for supporting the intended functionalities.

#### 4. Agents, Governance, and Decision-Making

Autonomous predictive analytics operating on a multi-cloud big data architecture require a system of agents that govern the architecture and drive decision-making according to a formal protocol. The principal agent oversees the overall design and operation of the architecture and also acts as a cloud-service integrator. The demand-predicting agent, which operates autonomously within the architecture, provides the credit-risk models that inform the credit-risk agents of the network of credit-risk model consumers. In their roles as consumers of the demand-predicting agent's model services, the credit-risk consumer agents dynamically request predicted demand and associated uncertainty information. Agents acting for inventory financiers use that

information to indicate through reconciliation signals for supply chain reconciliation and financing of expected order-fulfilment inventory.

Automated predictive analytics operating on a multi-cloud big data architecture require a network of agents capable of acting both independently and collaboratively. Governance ensures that all activities occur in accordance with the established design principles, whereas coordination establishes the protocols that agents of different types apply when communicating with each other. The principal agent oversees the architecture, integrates cloud-service offerings from different providers, and thus also plays a key role in guaranteeing data privacy and security.

#### Equation 2: Demand forecasting metrics mentioned

Assume we have:

- actual demand:  $y_1, y_2, \dots, y_T$
- predicted demand:  $\hat{y}_1, \hat{y}_2, \dots, \hat{y}_T$
- error at time  $t$ :  $e_t = y_t - \hat{y}_t$

##### 1) MAE (Mean Absolute Error)

**Goal:** average magnitude of errors (ignoring sign)

Step-by-step:

1. Compute each error:  $e_t = y_t - \hat{y}_t$
2. Take absolute error:  $|e_t|$
3. Average:

$$MAE = \frac{1}{T} \sum_{t=1}^T |y_t - \hat{y}_t|$$

##### 4.1. Agent Roles and Capabilities

In the proposed multi-cloud big data architecture, analytics agents perform the following roles: Demand Forecaster Agent, which produces predictive demand volumes that





minimize sum of squared errors and uncertain supply chain costs; Inventory Reconciliation Agent, which matches actual demand with supply to signal financing requirements; Credit Risk Scoring Agent, which rates supply chain stakeholders' creditworthiness using risk factors identified by the Credit Risk Factorization Model and corresponding scoring models; source system agents, which ingest, clean, and validate data; Country Data Agent, which acquires and serves country-level historical data sets to the Demand Forecaster Agent's Facebook-based deep-learning predictive models; Cloud Data Ingest Agent, which continuously acquires new country historical data and communicates to the Country Data Agent; Financial System Agent, which retrieves interest rate and exchange rate data from the Fed and communicates this information to the Credit Risk Scoring Agent; Data Bank Agent, which continuously monitors big data bank availability to ensure efficient time-series-data reuse; Data Provisioning Agent, which fulfills data processing and provisioning requests; and Operational Agents, which control operationalization of the Demand Forecasting Model, Inventory Reconciliation Model, and corresponding Financing Signal.

Each agent has access to all inputs and outputs of models developed by other agents and directly controls scheduling of its assigned models. For example, upon validation of the national-level predictive-demand Volumes-Error report, the Demand Forecaster Agent signals the Inventory-Reconciliation Agent to reconcile actual demand and supply and produce the Financing Signal report. Further justification for these capabilities lies in game-theoretic research showing that effective decision making is closely related to team membership and specific policy objectives.

#### 4.2. Coordination and Control Protocols

Patented multi-cloud data landscapes foster data interconnectivity and combined analytics across independent organizations. Security, privacy, and regulatory constraints inherently limit data sharing and disclosure of proprietary information. Within these environments, multi-instance agent-based control protocols and frameworks coordinate

data processing and analytical operations without revealing individual business intelligence. Cloud providers may leverage transaction information, credit risk forecasts, demand projections, inventory financing signals, and production indicators to autonomously allocate system tasks, harnessing their superior processing capabilities for predictive analytics.

The Composite Connect Procedure governs operation of non-communicating agents designed to share only information characterized as common knowledge. During deployment, agents request assistance from public cloud provider agents whenever they cannot cover a processing demand. When possible, such requests are satisfied by agents responsible for non-sensitive tasks. For a requested predict-and-query task to be eligible for external processing, demand-sensitive variables in the corresponding forecast must exceed a value determined by the querying agent. Composite Cloud Authorization Protocol automates the assignment of transactions with conflicting roles on non-production instances.


## 5. Predictive Demand and Inventory Optimization

Functioning at a national level, the proposed multi-cloud architecture supports predictive analytics frameworks for demand forecasting and inventory reconciliation in food service supply chains. Demand predictions for food service categories at state, region, and national levels are produced to support service distribution and supply chain operation using a combination of modeling, machine learning, and time series methods. Reconciling inventories between food service customers and their distributors, based on predictive demand signals, generates information for bank inventory financing decisions and bonding capabilities, while also aiding middle-mile distribution planning. Results from prediction experiments are compared based on root mean squared error (RMSE), mean absolute error (MAE), mean



absolute percentage error (MAPE), and mean bias error (MBE) metrics.

Current cloud-based supply chain forecasting and planning methods focus on the private partners within private cloud environments. However, in the food service sector of national economies, food service businesses rely on distributors such as broad-line food service distributors, that provide significantly higher inventories than needed for last-mile delivery to customer locations within the networks. These higher inventories provide the distributors the ability to offer higher service percentages to their customers than needed for maintaining low costs. At the same time, the food service businesses create demand surges such as during game days and holiday seasons which need to be accommodated with quick delivery times and minimal shortages. But these surges are small and infrequent and the higher inventories of the distributors are not required to maintain low costs. As a result, the difference between the distributor inventory and food service demand needs to be utilized for bonding and obtaining bank inventory financing for the distributors.



**Fig 3: AI-based predictive analytics for enhancing data-driven supply chain optimization**

## 5.1. Demand Forecasting Methods and Metrics

Proposed Agentic AI-Based Multi-Cloud Big Data Architecture Accuracy of prediction methods, generalization ability across time and space, and standardization of a set of

measurement metrics constitute crucial aspects of predictive demand forecasts in food service supply chains. The statistical methods proposed by Barlow and Prokop, Egri et al. and Xu et al., as well as ensemble approaches made up of these techniques, build on cohort-level time series histories and were validated against actual food service purchases by guests on stream. Using the methods, requirements were forecast for months  $N + 2$  to  $N + 5$ , to allow sufficient lead time for suppliers with salad-ingredient inventory to make growing-season decisions. Evaluation focused on the 2021-summer and 2021-fall seasons because the need for controlled foliage-vegetable product supply was especially high.

The two severest predictors of seasons with lower-than-expected purchases of salads are reduction in pandemic-related restrictions, and economy competitiveness; the the biggest and second-largest decreases are expected in months  $N + 4$  and  $N + 5$ , respectively. According to the validation, predicted purchases of salad-ingredient products, including other multi-ingredient salads, reached higher figures during the considered forecasting the period. Accuracy during the summer is high for salads and Body sampson; during the fall it increases or remains stable for all products except for other multi-ingredient salads. Simultaneously forecasts produce satisfying results, remaining within 10% of the real values in both seasons for both salads and other multi-ingredient salads. Forecasting uncertainty is negligible for salads over the summer and holds within 80% quantiles for Body sampson.

## 5.2. Inventory Reconciliation and Financing Signals

The demand signal is naturally complemented by an inventory signal, derived from the reconciliation. The joint calibration of demand and inventory ensures a proper match and provides a secondary signal for financing. Inventory data are undoubtedly sensitive, as they reveal the retailer's purchasing pattern, contingencies, and possibly even their marketing strategy. Although a food-temperature-related industry has seen the adoption of cloud-based supply chain solutions, general commerce lacks the equivalent.





Nevertheless, with global food distribution being in the spotlight due to food waste and other issues, a critical mass should expose such tendencies. The calibration algorithm outlined earlier should thus determine anomalous inventory levels. When the actual level is significantly above the predicted level, a stockout situation is eliminated. A deviation of more than, for example, five days from the predicted N stock level signifies an extremely low stocking level of the item: the eventual product failure can be avoided only by a special purchase.

These two inventory situations provide two financing signals that are contradictory in nature. In other words, the calibration suggests external financing only when needed, with no dead stock piling up. The flow-on effect is extended to other suppliers, stimulating production cycles and factory utilization or capacity buildup. A required setup can, for example, take care of hedging cost variability as well as ingredient manufacturing lead times. A strategic role is added to building blocks, with financing partners either providing this sourcing ability or being leveraged into such partner setups by leveraging the combined scale of suppliers, distributors, and retailers. When inventory is above the required level and getting old, a preferred change in stock quality—possibly analysed through a link to the Schneider matrix—is advised.

## 6. Credit Risk Modeling for Supply Chain Financing

In a capital-constrained environment, suppliers often do not have the capacity to finance inventory in response to higher incoming demand signals from customers. Suppliers with no credit history or who have exhibited past defaults face even higher barriers to obtaining financing as lenders would be unwilling to take on those credit risks or may charge prohibitively high interest rates. To facilitate financing for suppliers, a credit rating based score is an important component of the multi-cloud big data architecture and should signal lenders whether their investments are

warranted. The scoring model is thus an objective predictor of creditworthiness, developed using historical loan performance data and used on an ongoing basis during the loan screening process to ensure compliance with lending conditions.

However, the objective drivers of risk within the data may change over time and new high-risk drivers, not previously seen in the data, may emerge. The availability of historical bad data provides the opportunity to examine the credit risk of lenders in a data-driven manner, rather than it being solely driven by subjective gut-feeling and experience of the risk manager using a checklist approach, following which a credit decision is taken. A careful examination of bad data lends itself well to understanding how loans may go bad. A subsequent audit trail on loans, which have gone bad assists in determining which of the environmental factors is important for the loan to remain in the green zone.

### Equation 3: Inventory reconciliation + financing signals

Let:

- $D_t$  = actual demand rate (units/day) during period  $t$
- $\hat{D}_t$  = predicted demand rate (units/day)
- $I_t$  = actual on-hand inventory (units)
- $\hat{I}_t$  = predicted “required” inventory (units) implied by predicted demand and desired service level
- Convert inventory to **days-of-supply**:

$$\text{DoS}_t = \frac{I_t}{\hat{D}_t + \epsilon}, \widehat{\text{DoS}}_t = \frac{\hat{I}_t}{\hat{D}_t + \epsilon}$$

(where  $\epsilon$  is tiny, to avoid divide-by-zero)

Define deviation in days:

$$\Delta_t = \text{DoS}_t - \widehat{\text{DoS}}_t$$



Now encode the paper's "5 days" idea (example threshold):

- low-stock alarm if  $\Delta_t < -5\text{days}$  (inventory too low)

Agentic AI-Driven Multi-Cloud B...

- overstock/aging alarm if  $\Delta_t > +5\text{days}$  (inventory too high)

A simple **financing signal** could be:

$$S_t = \begin{cases} +1 & \Delta_t < -\tau(\text{finance urgent replenishment}) \\ 0 & |\Delta_t| \leq \tau(\text{no action}) \\ -1 & \Delta_t > \tau(\text{reduce / re-balance / avoid financing}) \end{cases}$$

with  $\tau = 5\text{days}$  as described in the paper's example

### 6.1. Risk Factors and Feature Engineering

To support demand shocks, food business operators often seek financing solutions based on trade credit and inventory financing. A well-functioning credit risk modeling system is therefore critical for lenders to maintain their profitability while reducing the risk of credit defaults. Credit assessments and predictions can be conducted by multiple institutions at different levels in the supply chain (e.g., suppliers, distributors, retailers). However, credit scoring models must be re-evaluated from time to time or at least monitored for their performance over time. Libraries from the credit scoring literature—such as those found in the R package *creditmodel*—provide various techniques and tools for developing and validating credit scoring models.


Credit risk factors at various levels in food supply chains typically include business characteristics (e.g., location, age) and industry performance indicators (e.g., inflation, interest rates, gross domestic product). The relationships between these risk factors and scores are sometimes non-linear. Feature engineering therefore also includes a filtering process based on a principal-component analysis of the risk factors and the use of a penalized regression approach

(LASSO) for building credit scoring models. Such effectiveness analysis can be guided by the R package *purrr*.

### 6.2. Scoring Models, Validation, and Monitoring

Multiple classification methods have been applied in predicting credit risk. To find the most effective algorithm, a comparative study has been conducted utilising six popular classification techniques such as logistic regression, support vector machines, K-nearest neighbours, Naive Bayes, random forest, and extreme gradient boosting. The scoring models were developed using 80% of the complete dataset, while the remaining 20% served as the test input. To evaluate the robustness and generalisation capabilities of these models, the 10-fold cross-validation technique was employed. Later on, the optimal models from these algorithms were selected for predicting the credit risk of a specific bank and companies listed within the bank's exploration area. Their performances are assessed based on the confusion matrix, area under receiving operator characteristic (ROC) curve, precision, recall, F-measure, accuracy, and Matthews correlation coefficient.

Feature importance results are calculated based on the random forest approach. The above six algorithms are employed for the prediction of credit risk and for optimal solution. Hyperband is applied. Hyperband is a hyper-parameter tuning technique that applies successively deeper random search on increasingly smaller resource allocations, enabling it to find good hyper-parameter configurations faster than random search and reduces the experimental budget for hyper-parameter search. Among them Support Vector Machine has been detected as the optimal algorithm in this scenario for predicting credit risk for banks. Risk scoring model is successfully developed and validated to predict the risk involved for loan applicants in banks.




## 7. Operationalization in National Food Service Networks

Deployment in Heterogeneous Cloud Environments comprises a Mobile Decision App in a National Cloud and Edge Mobile Nodes (EMN) in Regional Clouds. The EMN support Voice-Activated Demand Queries and Inventory Signals for Unconscious Decision-Making. At-Risk Suppliers and New Product Launches are Monitored from Public Clouds. Interoperability with Suppliers, Distributors, and Retailers is facilitated by Multi-Tenancy and Role-Based Access Controls in a National Cloud.

Multi-cloud Big Data Architecture deploying Agentic AI and Autonomous Analytics. A food service network is presented comprising Suppliers, Distributors, a Central Decision Maker and a Bank. Demand is forecasted independently in the cloud resources of the independent product suppliers using time series methods. However, the Demand Forecasting Module can also be run in the public cloud of the National Monitoring Agency that governs new product launches. The predicted demand is reconciled with the inventory levels of the Distributors to produce financing signals for the Bank and to generate alerts for at-risk

Suppliers. If the negative working capital of a Distributor exceeds a specified threshold, Automated Decision-Making Agents warn the Bank. The Bank maintains a scoring model to assess the credit risk of Suppliers.



**Fig 4: Application of Artificial Intelligence and Big Data in the Food Industry**

### 7.1. Deployment in Heterogeneous Cloud Environments

An implementation of the proposed multi-cloud big data architecture was developed within a national food service supply chain for the purpose of analyzing predictive demand, credit risk, and inventory financing. Agentic AI functionality was incorporated across the ingestion, processing, management, and predictive modeling of big data acquired from a range of external sources—government registries and social media sites—as well as from a consortium of suppliers, distributors, and operators connected through the cloud platform.

The heterogeneous cloud environment is maintained in the on-premise private clouds of the participating supplier, distribution, and retail organizations, while the demand analysis and predictive analytics are performed in a public cloud. Agentic AI analyses and proprietary data—reconciliation signals, predictive credit risk scores, and financing recommendations—are shared with the participants according to the defined privacy restrictions.





## 7.2. Interoperability with Suppliers, Distributors, and Retailers

Agency for predictive demand and credit risk analytics generates micro-level signals for food service supply chains, enabling suppliers, distributors, and retailers to meet their specific needs while mitigating risk. For example, a distributor makes supplier and lender selections based on product-specific credit scoring, lending coverage, terms, and conditions. Suppliers benefit from financing offers that reduce or eliminate margins lost to unfinanced inventory financing, or from financing proposals from distributors with optimal terms selected across all of its lenders. Retailers minimize out-of-stock incidents and maximize service levels with supplier financing offers that stretch days of inventories at a lower cost than expected loss of customers or revenue.

Cloud deployment covers Singapore and Japan, enabling regionally optimal placement of cloud analytics with multi-tenancy data governance. Supply chain partners are registered in the National Data Intermediary, supporting all elements of Agentic AI. Brokerage security is facilitated through the cloud operators offering KMS, making secure, authenticated, and permissioned use of the common brokers, credential storage, KMS, billing, and other services. The BigQuery SQL-based analytics deployed on Google Cloud Platform—BigQuery, Vertex AI, Cloud Functions, and GCS—and an AWS account running PostgreSQL, WordPress, Keycloak, and Lambda functions operate are State and Territory solutions, ready to be deployed elsewhere using the multi-cloud data connectors.

## 8. Conclusion

Data privacy, confidentiality, and security are critical concerns for any data management operation, especially for big data applications in multi-cloud environments with the presence of multiple heterogeneous data owners and users. A proper data privacy and security framework is, therefore, necessary to protect sensitive data, restrict access for unauthorized entities, prevent data misuse, and control

queries posed by data consumers. Such a framework is equally essential to minimize the risk of tangible data leaks or irreversible compromise of privacy during the process. Furthermore, appropriate regulatory compliance requirements in food systems and, more specifically, food supply chains must be respected, to avoid any potential economic, health, and environmental risks. Confidential data about a partner in a supply chain must not be sent or even made available to any other partner out of this chain. Thus, regulatory compliance and data privacy must be managed consistently in such systems.

A cloud-based regulatory compliance framework for food systems is essential to create the preconditions for a sustainable food production at the global level. A machine learning-based regulatory compliance framework addresses regulatory compliance and data privacy, controlling the transmission of sensitive data in a declarative manner by applying different cloud-based techniques. In addition, a machine learning-based regulatory compliance framework fitted with appropriate modules can also associate multi-cloud service providers for any kind of big data analytics—including of edge-computing-based systems—by preserving the data privacy, confidentiality, and preventing data misuse.

### 8.1. Data Privacy and Security Frameworks

Public acceptance of predictive demand forecasting and supply chain financing hinges on equitable, transparent, and protected data sharing among the cloud services. Personal privacy protection is paramount to ensure citizens' readiness to use mobile applications that share health information with other sectors to optimize logistics. Despite the commercial interest in offering free information, any uncontrolled dissemination of private data opens serious ethical issues. Therefore, a biometric privacy-preserving data collector is needed, designed for data security, communication, and trusted processing, allowing the collection of sensitive biometrics in encrypted form. Data are aggregated in a cloud and a bioconversion module is used to convert the encrypted data sets into an optimally compressed form for mobile machine learning applications.



Deep learning translation techniques for fine-grained human activity recognition ensure that important behaviours can be captured even when they are performed in the presence of others or when two persons are close together. The models trained and tested in a local context are then deployed in a cloud-in-the-loop configuration, in which an external work cloud takes care of creating the cloud model update that is then sent back to overwrite the local model. Future integration with the COVID-19 contact tracing application is foreseen. Finally, a mixed architecture framework coordinates devices with lesser resources through private edge servers for the detection of particular exigent cases. On-device training allows a light transfer of model information without disclosing original models.

## 8.2. Regulatory Compliance in Food Systems

With food supply chains recognized as critical infrastructures around the world, the ability of food service establishments to remain open and operating also has public safety implications that go beyond the upkeep of business revenues and profitability. Monitoring and controlling credit risk across the entire service supply chain will, therefore, enable better allocation of community resources in support of healthy demand and service availability. Risk modeling and scoring components presented previously in the methodology can be developed within the context of a broader governance framework that addresses the unique requirements of national food systems. Such systems are not all-closely regulated but rather rely upon strategic controls imposed at various levels of the supply chain. Since food service establishments are the last link in food service supply chains, the assessment of their creditworthiness must take account of the health and safety compliance ratings of their immediate suppliers. The frequency of inspection with official grading of compliance acts as a temporary constraint on any hospitality or food service offering. Provisioning for services with liquor licenses is similarly governed by the degree of risk associated with offering alcoholic beverages. Within the North American context, additional municipal approvals are required to set-up temporary or permanent

booths or kiosks if the food service facilities is not "business as usual."

Regulatory and governance presence therefore influences network design, provisioning stability, continuity during retail seasons, and tangible ratings from trusted sources. Safe browsing and payment services offered by short-term leasing platforms, on the other hand, reduce friction for consumers, while their accredited suppliers add levels of trust and confidence in secure browsing and purchasing. Other suppliers and service partners depend on word-of-mouth recommendations. These factors can be captured in the credit risk signals available for consumption and non-consumption lending in food service supply chains and networks.

## 9. References

- [1] Akhtar, P., Tse, Y. K., Khan, Z., & Rao-Nicholson, R. (2019). Data-driven and adaptive leadership contributing to sustainability and competitive advantage: The role of big data analytics capabilities. *Journal of Business Research*, 101, 568–578.
- [2] Amazon Web Services. (2023). Well-architected framework. Amazon Web Services.
- [3] Armbrust, M., Das, T., Sun, L., Yavuz, B., Zhu, S., Murthy, M., Torres, J., van Hovell, A., Wang, A., Ionescu, M., & Xin, R. (2020). Delta Lake: High-performance ACID table storage over cloud object stores. *Proceedings of the VLDB Endowment*, 13(12), 3411–3424.



- [4] Benjelloun, F. Z., Lahcen, A. A., & Belfkih, S. (2021). Big data analytics architectures: A survey. *Journal of King Saud University – Computer and Information Sciences*, 33(3), 257–272.
- [5] Berg, T., Saunders, A., Steffen, S., & Streitz, D. (2020). Mind the gap: The difference between U.S. and European loan spreads. *Review of Financial Studies*, 33(2), 948–987.
- [6] Bertsimas, D., Kallus, N., Weinstein, A. M., & Zhuo, Y. D. (2020). Retail forecasting: An optimal machine learning approach. *Management Science*, 66(7), 3013–3033.
- [7] Bhamra, R., Dani, S., & Burnard, K. (2021). Resilience: The concept, a literature review and future directions. *International Journal of Production Research*, 59(6), 1521–1540.
- [9] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- [10] Chopra, S., & Meindl, P. (2022). *Supply chain management: Strategy, planning, and operation* (8th ed.). Pearson.
- [11] Chen, H., Chiang, R. H. L., & Storey, V. C. (2019). Business intelligence and analytics: From big data to big impact. *MIS Quarterly*, 43(4), 1165–1188.
- [12] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- [13] Cohen, M. A., & Kouvelis, P. (2021). Supply chain finance: Research, practice, and the opportunities for analytics. *Foundations and Trends in Technology, Information and Operations Management*, 14(1), 1–98.
- [14] Dolgui, A., Ivanov, D., & Sokolov, B. (2020). Reconfigurable supply chain: The X-network. *International Journal of Production Research*, 58(13), 4138–4163.
- [15] Dunning, T., & Friedman, E. (2019). *Introduction to Apache Spark: A unified analytics engine for big data*. O'Reilly Media.
- [16] Fettke, P., Loos, P., & Zwicker, J. (2021). Reinforcement learning in operations management: A systematic literature review. *European Journal of Operational Research*, 292(1), 1–17.
- [17] Fisher, M. L. (1997). What is the right supply chain for your product? *Harvard Business Review*, 75(2), 105–116.
- [18] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.





- [19] Gürtler, M., & Eickhoff, M. (2016). Supply chain finance, optimization, and the allocation of liquidity. *Journal of Business Economics*, 86(1), 1–22.
- [20] He, H., Zhang, W., & Zhang, S. (2021). A multi-cloud computing framework for scalable big data analytics. *Future Generation Computer Systems*, 120, 27–40.
- [21] Ivanov, D., & Dolgui, A. (2020). Viability of intertwined supply networks: Extending the supply chain resilience angles towards survivability. *International Journal of Production Research*, 58(10), 2904–2915.
- [22] Jüttner, U., & Maklan, S. (2019). Supply chain resilience in the global financial crisis: An empirical study. *Supply Chain Management: An International Journal*, 24(6), 677–690.
- [23] Keerthi, S. S., & Lin, C. J. (2003). Asymptotic behaviors of support vector machines with Gaussian kernel. *Neural Computation*, 15(7), 1667–1689.
- [24] Kouvelis, P., Dong, L., Boyabatlı, O., & Li, R. (2021). Introduction to the special issue on supply chain finance. *Production and Operations Management*, 30(2), 301–303.
- [25] Larocque, P., & Lakhani, K. R. (2020). The emergence of data-driven supply chains. *California Management Review*, 62(3), 5–19.
- [26] Lim, M. K., Mak, H. Y., & Rong, Y. (2022). Toward explainable AI in supply chain analytics: A review and research agenda. *International Journal of Production Research*, 60(22), 6990–7011.
- [27] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*, 36(1), 54–74.
- [28] Mirzaei, M., Naderi, B., & Tavana, M. (2023). Machine learning models for credit risk assessment: A comparative study. *Expert Systems with Applications*, 213, 118918.
- [29] Mishra, D., Gunasekaran, A., Papadopoulos, T., & Dubey, R. (2019). Supply chain performance measures and metrics: A bibliometric study. *International Journal of Production Research*, 57(7), 1–21.
- [30] Rahman, M. H., & Subramanian, N. (2024). Predictive analytics for demand forecasting in food supply chains: A systematic review. *Computers & Industrial Engineering*, 187, 109873.
- [31] Zhou, L., Pan, S., Wang, J., & Vasilakos, A. V. (2019). Machine learning on big data: Opportunities and challenges. *Neurocomputing*, 331, 1–5.

