© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723457

AI-Assisted Demand Forecasting and Production Scheduling for High-Volume Food Manufacturing in Saudi Arabia: A Food Security Framework for Vision 2030

ATIF MAQBOOL SYED

Abstract- High-volume food manufacturing in Saudi Arabia is increasingly exposed to volatile demand, short product life cycles, promotion-driven peaks, imported input dependencies, capacity bottlenecks, and the operational consequences of seasonal and religious events. These conditions make demand forecasting and production scheduling inseparable rather than sequential planning tasks. This review synthesizes research published between 2020 and 2025 on artificial intelligence-assisted demand forecasting, perishable-product planning, finite-capacity scheduling, digital-twin-enabled decision support, and food supply-chain resilience. The study develops an integrated framework in which demand sensing produces probabilistic forecasts that are translated into production decisions through capacity, shelf-life, sanitation, changeover, labour, inventory, and service constraints. Evidence indicates that machine-learning and deep-learning methods can improve forecast accuracy when promotions, weather, calendar events, regional variation, and censored demand are represented appropriately, while optimization models can convert these signals into feasible schedules that reduce changeovers, waste, shortages, and unstable resource use [1-8]. The review also identifies a persistent implementation gap: forecasting studies often stop at predictive accuracy, whereas scheduling studies frequently assume demand inputs are fixed and reliable. The proposed framework closes this gap by linking forecast uncertainty, rolling-horizon scheduling, execution feedback, and food-security outcomes. For Saudi manufacturers, the value of this integration lies not only in lower cost, but also in improving product availability, freshness, localized processing capacity, shock recovery, and evidence-based management under Vision 2030. The paper concludes with a staged implementation roadmap and research priorities for explainability, cross-site scheduling, event-aware forecasting, data governance, and resilient planning.

I. INTRODUCTION

Saudi food manufacturers operate in a planning environment where demand volatility, perishability, and capacity scarcity interact continuously. A forecast that is statistically accurate but operationally unusable can still generate shortages, expiry, overtime, excessive changeovers, or poor service. Conversely, an optimized production schedule based on a weak or biased demand signal may allocate scarce lines and labour efficiently to the wrong product mix. Recent research therefore points toward tighter integration between predictive analytics and operational optimization rather than treating forecasting and scheduling as independent functions [1-4]. This is especially relevant in high-volume plants producing dairy, beverages, bakery items, processed meat, frozen foods, confectionery, and packaged staples, where thousands of stock-keeping-unit and location combinations may need daily or weekly planning.

Artificial intelligence has widened the range of demand signals that can be processed. Beyond historical sales, contemporary models can incorporate promotions, weather, macroeconomic variables, regional patterns, online activity, special days, and demand censoring [5-10]. Deep-learning and ensemble approaches can capture nonlinear interactions and temporal effects that conventional time-series models may miss, but the operational value of accuracy depends on calibration, stability, interpretability, and the way forecast uncertainty is transferred into downstream decisions. In perishable categories, forecast error is asymmetric: over-forecasting increases expiry and markdown risk, whereas under-forecasting may create lost sales, production instability, or emergency replenishment. This asymmetry means that planners need probability distributions, quantiles, or scenario ranges, not only single-point forecasts.

Production scheduling in food processing is equally distinctive. Plants often combine batch and continuous stages, shared utilities, allergen restrictions, cleaning requirements, sequence-dependent setups, labour calendars, packaging constraints, minimum run lengths, and microbiological or temperature requirements. Optimization research demonstrates that mixed-integer programming, decomposition, rolling-horizon methods, and integrated production-distribution models can substantially improve industrial schedules when these constraints are represented explicitly [11-16]. Digital-twin approaches extend this logic by connecting planned schedules to live operational states and enabling what-if analysis before changes are released [17,18].

The Saudi context strengthens the strategic importance of this integration. National industrial policy emphasizes food processing, localization, resilient supply chains, and higher domestic capacity as part of broader economic diversification and food-security objectives [25-30]. High-volume manufacturing can contribute to those goals only when local plants can anticipate demand, absorb disruptions, use capacity efficiently, and protect service levels during peak periods. The central argument of this review is therefore that demand forecasting should be judged by its ability to improve production decisions, while scheduling should be designed to react explicitly to uncertainty rather than assuming a single deterministic demand input.

II. AIM, OBJECTIVES AND REVIEW QUESTIONS

The aim of this review is to develop a decision-support framework that connects AI-assisted demand forecasting with production scheduling for high-volume food manufacturing in Saudi Arabia, while explicitly linking operational performance to food-security outcomes. The focus is not on proposing one universally superior algorithm. Instead, the study seeks to explain how forecast generation, uncertainty representation, capacity-aware scheduling, execution feedback, and management governance can operate as one planning system.

Five objectives guide the review. First, it evaluates recent forecasting approaches and identifies where machine learning, deep learning, probabilistic forecasting, and demand sensing create practical value for food-related demand patterns [5-10]. Second, it examines scheduling methods that address perishability, multi-product lines, changeovers, resource constraints, routing, and rolling replanning [11-16]. Third, it synthesizes evidence on the interface between prediction and optimization, including digital twins and dynamic decision support [17-20]. Fourth, it translates the literature into a Saudi manufacturing architecture that reflects event-driven demand, localization priorities, imported-input risk, and the need for reliable domestic supply [25-30]. Fifth, it identifies governance and implementation conditions required to prevent advanced analytics from becoming disconnected pilot projects.

The review addresses four questions. What types of demand signals and modelling approaches are most suitable for volatile, high-volume food categories? How should uncertainty be transferred from forecasts into feasible production schedules? Which operational and food-security measures should be used to judge the combined system? Finally, what implementation sequence can allow Saudi manufacturers to move from fragmented spreadsheet planning toward governed, adaptive, and scalable decision support?

III. REVIEW METHODOLOGY

A structured integrative review method was adopted because the research problem spans forecasting, operations research, food manufacturing, digital decision support, and supply-chain resilience. The review window was restricted to 2020-2025 to reflect current analytical capabilities and operational conditions. Searches were organized around four concept groups: artificial intelligence and demand forecasting; food or perishable-product production planning; finite-capacity or dynamic scheduling; and Saudi food security, industrial localization, or supply resilience. Priority was given to peer-reviewed journal articles, conference papers with clear methodological contribution, and official Saudi or international institutional reports relevant to food-security planning. Studies were included when they contributed directly to at least one element of the proposed decision chain: demand sensing, forecast modelling, uncertainty representation, production or distribution scheduling, digital-twin support, waste reduction, or resilience. Studies were excluded when their main contribution concerned unrelated computer-vision quality inspection, consumer recommendation, purely agricultural yield prediction, or logistics problems without a meaningful planning connection. The evidence was then coded by decision level, data requirements, method, operational constraint, performance measure, and practical implication. This thematic coding approach was selected because the reviewed studies are methodologically heterogeneous; direct meta-analysis would be inappropriate given different datasets, industries, horizons, loss functions, and objective formulations.

The synthesis deliberately distinguishes predictive performance from operational performance. Forecasting papers were interpreted not only through error measures, but also through their relevance to bias, promotions, special days, freshness, and demand uncertainty [5-10]. Scheduling papers were interpreted through feasibility, cost, waste, service, changeovers, inventory, energy use, and responsiveness [11-16]. Digital-twin and resilience studies were used to assess how real-time information and scenario modelling can connect planning with execution [17-24]. Finally, Saudi strategic documents were analyzed for their implications for domestic processing capacity, local supply, resilience, and industrial competitiveness [25-30]. The resulting framework is therefore a conceptual synthesis grounded in recent evidence rather than an empirical test of one plant or dataset.

IV. AI-ASSISTED DEMAND FORECASTING IN FOOD MANUFACTURING

Food demand has structural features that make model selection consequential. Sales can be affected by weekly seasonality, promotions, weather, pay cycles, holidays, religious periods, tourist flows, regional differences, product substitution, stockouts, and rapid changes in consumer behaviour. Predictive models that use only lagged sales may therefore underperform when exogenous factors explain a meaningful share of variation. Punia and Shankar demonstrate the value of combining structured and unstructured variables in an ensemble framework for packaged-food demand, including promotions, weather, economic conditions, and media signals [5]. Broader reviews similarly show that machine-learning and deep-learning methods are increasingly used when relationships are nonlinear or data volumes are large [6,7].

For Saudi manufacturers, event-aware forecasting is particularly important because Ramadan, Eid periods, Hajj and Umrah flows, school calendars, weather extremes, and regional consumption patterns can alter both total volume and product mix. These events should not simply be encoded as binary holiday flags. Their effects may begin before the event, vary by region and channel, and interact with promotions or distribution stocking. A practical forecasting architecture should therefore separate baseline demand from event uplift and allow planners to compare expected, upside, and downside scenarios. Hierarchical reconciliation is also valuable because forecasts may be required simultaneously at national, region, customer, plant, product-family, and SKU levels. Coherence across these levels matters when production decisions are made at a more aggregated level than customer orders.

Food demand is also frequently censored. A low recorded sale may reflect low demand, but it may also reflect an out-of-stock situation. Miguéis et al. show that accounting for censored demand can improve forecasting for fresh fish while supporting both availability and waste reduction [8]. Related retail work demonstrates that machine-learning models can handle large-scale daily forecasting and special-day effects more effectively than simpler alternatives in many contexts [9]. Bayesian approaches provide another useful perspective because they express uncertainty explicitly and can improve food-order decisions where waste and availability are jointly important [10].

No single algorithm should be selected for all SKUs. Stable high-volume products may be forecast adequately with exponential smoothing or autoregressive baselines, whereas promotion-sensitive or highly nonlinear series may justify gradient boosting, recurrent neural networks, convolutional-recurrent hybrids, or ensembles. Model portfolios reduce dependency on one technique and allow automatic selection according to history length, intermittency, volatility, and business value. However, the governance layer is as important as the model layer. Planners require consistent definitions of demand, clear treatment of returns and stockouts, versioned event calendars, bias monitoring, and a documented override process. Human overrides should be retained when commercial information is genuinely unavailable to the model, but every override should be measured for forecast value added. This prevents intuition from becoming an unchallenged source of systematic bias.

Accuracy measures should also match business consequences. Mean absolute percentage error can become unstable for low-volume items, while root mean squared error over-penalizes large deviations. Weighted absolute percentage error, bias, pinball loss for quantile forecasts, and service-weighted measures are often more informative for high-volume food portfolios. Most importantly, forecast quality should be assessed downstream. A model that lowers statistical error but increases schedule nervousness, expiry, or emergency overtime may not create operational value. This requirement leads directly to the scheduling interface.

V. PRODUCTION SCHEDULING UNDER FOOD-MANUFACTURING CONSTRAINTS

Production scheduling converts forecast information into decisions about what to make, where, when, and in what sequence. Food plants complicate this task because products are perishable and many process constraints are sequence dependent. Georgiadis et al. demonstrate how mixed-integer programming can address large-scale food-processing schedules with continuous and batch operations while respecting industrial constraints [11]. Angizeh et al. similarly show that multi-line scheduling can jointly consider productivity, scrap, energy, changeover duration, and labour cost [12]. These studies illustrate why a forecast cannot be translated directly into planned quantities without considering the technical capability of the factory.

Rolling-horizon methods are particularly useful when demand and operating conditions change after the weekly plan is created. Kamhuber et al. show how production smoothing can balance multiple objectives under dynamic capacity and stock targets in a food-manufacturing setting [13]. Closed-loop scheduling research in canned-food production further demonstrates the importance of connecting control and scheduling when shared resources or process states change [14]. This logic is essential for Saudi plants facing short lead times, temperature-sensitive operations, or imported packaging materials whose availability may change during the horizon.

Perishability changes inventory policy. Conventional safety-stock logic can encourage excess production, but products with limited shelf life cannot be buffered indefinitely. Polotski et al. demonstrate production-control policies in which inventory targets adapt to periodic demand while being bounded by shelf-life considerations [15]. Production-routing research extends the decision problem beyond the factory by considering age, freshness, inventory, and distribution together [16]. These findings imply that production schedules should include remaining shelf life, destination lead time, and expected consumption speed rather than treating all finished-goods inventory as equivalent.

Changeover logic is another critical issue. High-volume plants often face flavour, allergen, package-size, or sanitation sequences that create asymmetric setup times. A schedule that minimizes total run time but ignores cleaning matrices may be infeasible or unsafe. Likewise, minimum run lengths and packaging-line constraints can force production above immediate demand, requiring the forecast layer to communicate uncertainty ranges so that planners understand the risk of overproduction. Multi-objective scheduling is therefore more realistic than pure cost minimization. Relevant objectives include service level, expiry, schedule stability, overtime, changeovers, energy, labour utilization, and inventory age.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723457

The best scheduling method also depends on decision speed. Exact mixed-integer optimization is appropriate when the model can solve within the planning cycle, but large plants may require decomposition, heuristics, or hybrid methods. Recent reviews of machine learning for scheduling highlight the growing role of reinforcement learning and hybrid AI-optimization approaches in dynamic environments, while also warning about data quality, computational complexity, and generalization [21]. For practical adoption, manufacturers should therefore use the simplest method that captures the binding constraints and produces repeatable decisions within the required response time.

VI. INTEGRATING FORECASTING AND SCHEDULING

The central design problem is the handoff between prediction and optimization. Many forecasting systems produce one expected number per SKU-period, while many scheduling models treat that number as certain. This creates false precision. An integrated system should instead transfer both expected demand and uncertainty. Quantile forecasts, prediction intervals, or scenario sets can be converted into service targets, safety requirements, flexible production quantities, or reserve-capacity rules. The scheduling engine can then evaluate the cost of protecting against upside demand relative to the waste risk of overproduction.

A practical architecture begins with demand sensing, where sales, orders, promotions, event calendars, weather, stockouts, customer commitments, and regional signals are aligned at a common time grain. The forecasting layer generates a baseline, event uplift, probability range, and bias diagnostics. The decision layer then translates those signals into required inventory positions while accounting for existing stock, expiry dates, inbound materials, and service targets. Finite-capacity scheduling converts the requirement into line assignments, campaign sizes, sequences, labour demand, sanitation windows, and planned completion times. Execution systems return actual yield, downtime, changeover duration, production quantity, inventory age, and service outcomes to the analytical layer.
Digital-twin research provides a useful mechanism for closing this loop. Koulouris et al. describe how process and digital-twin models can support simulation and scheduling in food production [17]. Maheshwari et al. extend the concept toward real-time planning, monitoring, and control across food supply chains [18]. The value is not the visual replica itself, but the ability to test schedule changes against current process states before implementation. When a filler line fails, an ingredient arrives late, or demand surges unexpectedly, the digital model can compare alternative sequences and quantify the effect on service, waste, and capacity.

The integrated framework should also protect schedule stability. Constantly reacting to every forecast update can create nervousness, labour disruption, and material inefficiency. Frozen zones, change thresholds, and decision hierarchies are therefore necessary. Minor forecast changes may alter inventory projections without changing the production sequence, whereas larger deviations can trigger rescheduling. The threshold should depend on product criticality, remaining shelf life, capacity tightness, and the cost of change. In this way, AI supports disciplined decision-making rather than continuous automated turbulence.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723457

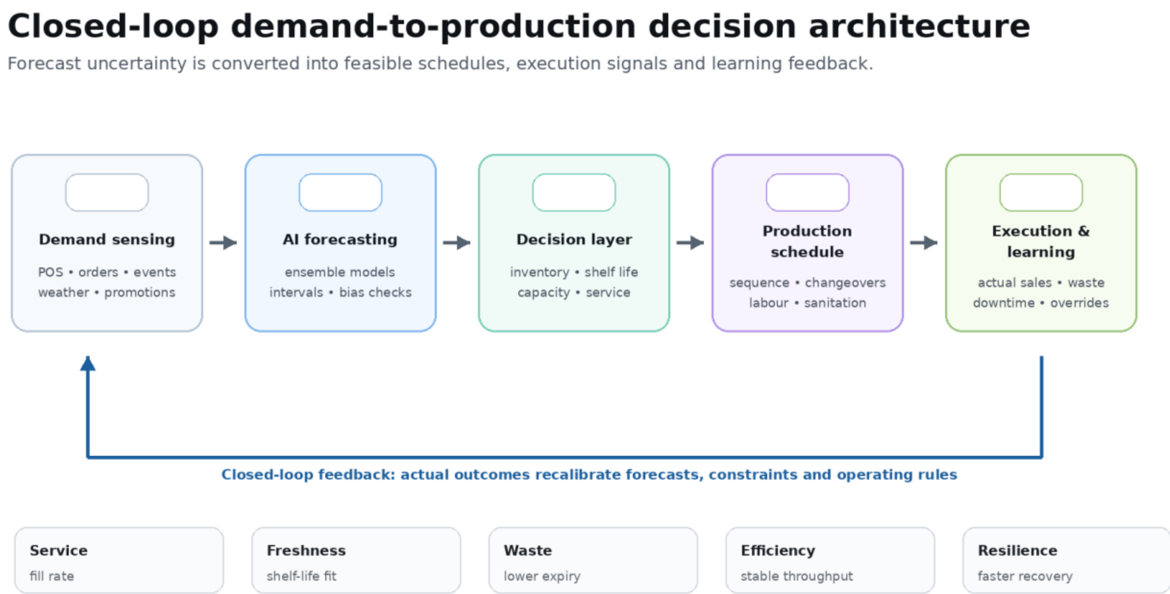

Figure 1. Integrated closed-loop architecture linking demand sensing, forecasting, scheduling, execution, and learning.

Table 1. Evidence synthesis for forecasting and scheduling methods in high-volume food manufacturing.

Decision area	Typical methods	Main contribution	Key limitation	Operational use
Baseline forecasting	ETS, ARIMA, seasonal naive	Transparent benchmark and stable seasonal patterns	Weak nonlinear and event interactions	Reference forecast and fallback
Machine learning	Gradient boosting, random forest, SVR	Captures nonlinear exogenous drivers	Requires engineered features and governance	Promotion and event-sensitive SKUs
Deep learning	LSTM, CNN-LSTM, ensembles	Learns long temporal dependencies	Data intensity and explainability	Large volatile portfolios
Probabilistic forecasting	Quantiles, Bayesian models, scenarios	Represents demand uncertainty explicitly	Calibration must be monitored	Service and reserve-capacity decisions
Finite-capacity scheduling	MILP, decomposition, rolling horizon	Creates feasible line sequences under constraints	Model complexity and solve time	Weekly and daily production plans
Dynamic rescheduling	Heuristics, hybrid AI-optimization	Fast reaction to disruption	May sacrifice optimality	Short-horizon recovery
Digital twin support	Simulation and live state models	Tests what-if scenarios before release	Integration and model-fidelity demands	Execution feedback and stress testing

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723457

VII. SAUDI FOOD-SECURITY AND VISION 2030 ALIGNMENT

Food security in a manufacturing context extends beyond agricultural production. It includes the ability to process, package, store, and distribute sufficient food reliably, particularly when external shocks affect imported materials, shipping capacity, or regional demand. Saudi industrial strategy places food processing among the sectors where localization, cluster development, and stronger supply chains can improve national resilience [25-27]. This creates a direct link between factory planning quality and strategic outcomes.

Integrated forecasting and scheduling can support availability by improving the probability that critical products are produced before demand peaks. It can support affordability by reducing waste, emergency production, inefficient changeovers, and excess inventory. It can support localization by helping domestic plants use installed capacity more effectively and demonstrate reliable service compared with imported alternatives. It can also support resilience by providing scenario-based responses when materials, logistics, or equipment become constrained. These contributions are especially relevant when a plant serves multiple regions or channels and must allocate limited capacity during a disruption.

Saudi food-manufacturing clusters may also benefit from shared planning signals. Forecast uncertainty at one manufacturer can propagate upstream to packaging, ingredients, cold storage, and transport. Where commercial arrangements permit, structured information sharing can improve capacity reservation and reduce the bullwhip effect. National statistical and food-security information can provide contextual indicators, while company-level systems retain the detailed transactional data required for execution [28-30]. The strategic opportunity is therefore not a centralized forecast for all firms, but interoperable decision practices in which data quality, event calendars, risk signals, and capacity assumptions are governed consistently.

The framework should be evaluated through a balanced set of measures. Operational metrics include forecast bias, weighted forecast error, line utilization, changeover hours, schedule adherence, waste, inventory age, overtime, and service level. Food-security-oriented metrics include critical-SKU availability, recovery time after disruption, domestic capacity utilization, dependency on emergency imports, and the proportion of demand met within freshness targets. Management should avoid optimizing one metric in isolation. Raising line utilization by producing long campaigns may reduce changeovers but increase expiry; maximizing freshness may increase short runs and cost. The decision-support system must expose these trade-offs rather than hide them behind a single score.

VIII. IMPLEMENTATION ROADMAP AND GOVERNANCE

Implementation should proceed in stages because forecast sophistication cannot compensate for weak master data or unclear planning ownership. The first stage is a trusted data foundation. Manufacturers should standardize SKU, customer, location, recipe, line, capacity, calendar, shelf-life, and promotion data; distinguish orders from true demand; and reconcile inventory and production histories. The second stage is forecast discipline. A model portfolio should be established with transparent benchmarks, event features, bias monitoring, and automated back-testing. Planners should receive both point forecasts and uncertainty ranges, while manual overrides are logged and evaluated.

The third stage is scheduling integration. Capacity, sanitation, changeover, labour, packaging, and material constraints should be codified in a finite-capacity model. Initial deployment can focus on one plant or bottleneck line where the economic value of better sequencing is measurable. The fourth stage is adaptive control. Forecast updates, material risks, downtime, and actual yields are fed into a rolling-horizon process with clearly defined rescheduling thresholds. The fifth stage is network resilience, where multiple plants, suppliers, and distribution regions are stress-tested using scenarios. This sequence prevents a common failure mode in analytics programmes: building sophisticated dashboards or algorithms before operational definitions and decision rights are stable.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723457

Governance requires cross-functional ownership. Commercial teams own promotion and customer intelligence; supply-chain planning owns demand consensus and inventory policy; manufacturing owns feasible capacity and process constraints; finance validates cost assumptions; information-technology teams manage data pipelines, security, and model deployment. A planning council should approve KPI definitions, change thresholds, and exception rules. Model performance should be reviewed separately from business performance so that a technically strong model is not retained if it creates poor operational outcomes.

Explainability also matters. Managers do not need every mathematical detail, but they need to understand why a forecast changed, which constraint caused a schedule decision, and what trade-off is created by an override. Scenario comparison is therefore more useful than a black-box recommendation. A planner should be able to see, for example, that protecting a ninety-fifth percentile demand scenario for a short-life product requires additional overtime and creates a quantified expiry risk if the uplift does not materialize. Decision logs then create organizational learning and provide evidence for refining future rules.

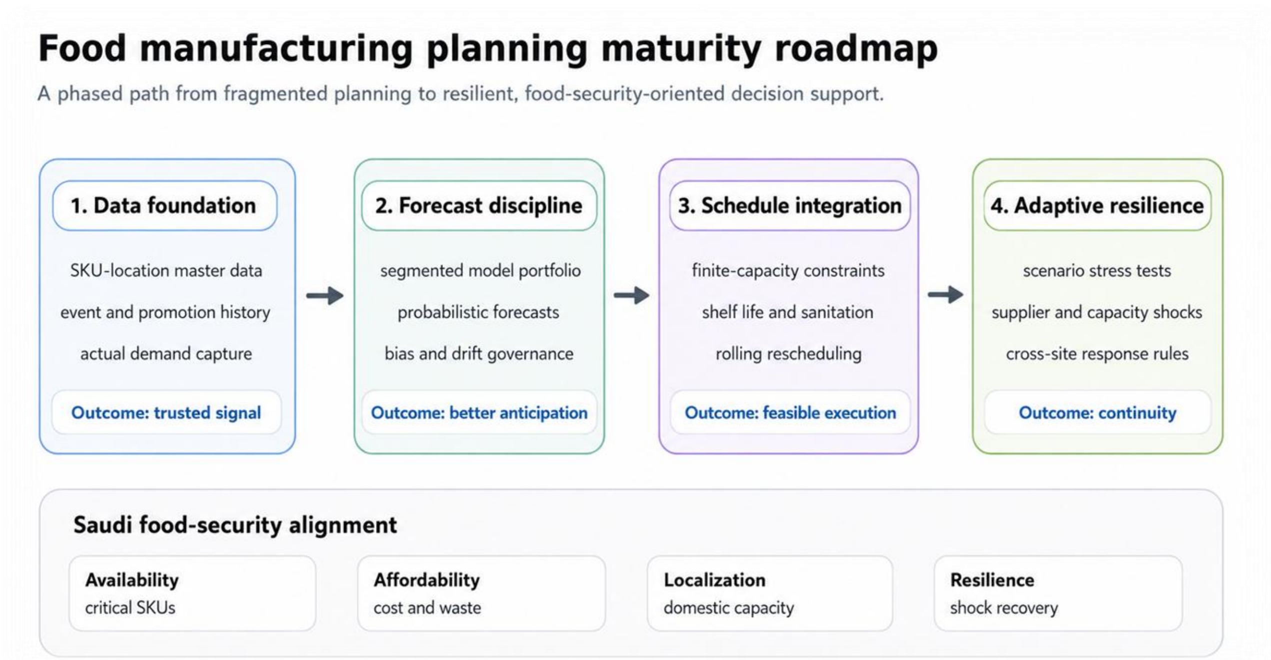

Figure 2. Maturity roadmap for moving from fragmented planning to adaptive and resilient food-manufacturing decision support.

Table 2. Decision-support framework linking data, planning decisions, KPIs, and food-security outcomes.

Framework layer	Priority inputs	Primary decision	Core KPIs	Food-security effect
Data foundation	Master data, sales, inventory, capacity, shelf life	Establish trusted planning state	Completeness, reconciliation, latency	Reliable visibility
Forecast layer	Events, promotions, weather, channels, regions	Generate baseline and uncertainty	Bias, WAPE, pinball loss, value added	Earlier demand anticipation
Inventory decision	On-hand age, inbound supply, service targets	Translate forecast into requirements	Freshness, stock risk, coverage	Availability with lower expiry

Schedule layer	Line rates, setups, labour, sanitation, packaging	Create feasible production sequence	Adherence, changeovers, overtime	Stable domestic conversion capacity
Adaptive control	Actual demand, yield, downtime, supplier risk	Trigger rolling reschedule when material	Service, waste, recovery time	Shock absorption
Network resilience	Multi-site capacity, critical materials, scenarios	Allocate capacity under disruption	Critical-SKU availability, recovery	Continuity and localization

IX. DISCUSSION, RESEARCH GAPS AND FUTURE DIRECTIONS

The reviewed evidence supports a shift from isolated predictive models toward integrated planning systems, but several gaps remain. First, there is limited research that evaluates forecasting models by their downstream scheduling consequences. Most studies optimize statistical error, while production models assume demand is given. Future research should compare forecasting methods according to total operational cost, waste, service, schedule stability, and resilience. Such work would identify cases where a slightly less accurate but better-calibrated forecast produces superior manufacturing outcomes.

Second, event-rich Saudi demand requires more contextual research. Religious periods, tourism flows, extreme temperatures, regional preferences, and public-sector procurement may interact differently across categories. Local datasets could support transfer-learning or hierarchical models that preserve information across similar products while adapting to regional behaviour. Third, cross-site scheduling is underdeveloped in food settings. Manufacturers with multiple plants need methods that decide not only when to produce, but also where, while accounting for freshness, logistics, local capacity, and disruption risk. Fourth, explainable and responsible AI needs deeper operational treatment. Explainability should extend beyond feature importance to decision consequences: which forecast driver changed the schedule, what capacity constraint became binding, and how much service or waste risk is associated with the recommendation. Recent research on explainable AI for perishable supply chains shows the potential of combining segmentation and transparent prediction [19]. Related work on machine-learning-based food-waste reduction demonstrates the environmental value of improved demand estimates [20]. Future systems should therefore integrate economic, environmental, and service objectives rather than optimizing forecast error alone.

Fifth, resilience should be modeled as a dynamic capability rather than a static safety stock. AI-enabled surveillance, digital twins, and scenario planning can help organizations detect weak signals and react faster [22-24]. However, additional redundancy is not always economically justified. Research should examine when flexible capacity, supplier diversification, strategic stock, or faster rescheduling provides the best protection. For Saudi manufacturers, this question is closely linked to localization: domestic capability creates strategic value when it can respond reliably under stress, not simply when nominal installed capacity exists.

A sixth gap concerns the relationship between planning frequency and organizational capacity. More frequent forecast refreshes do not automatically improve decisions if production teams cannot absorb continual change. Research should therefore investigate the economic value of different replanning cadences, frozen-zone lengths, and trigger thresholds across product families with different shelf lives and service priorities. This would help distinguish useful responsiveness from schedule nervousness. It would also clarify whether fast-moving products should be governed by hourly signals, daily decision cycles, or more stable campaign-based planning windows.

A seventh opportunity lies in integrating commercial behavior with physical production constraints. Promotions are typically evaluated through sales uplift, yet a promotion can create hidden

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723457

manufacturing costs through additional setups, packaging switches, overtime, and leftover inventory. Future studies should quantify promotion profitability after considering these operational effects. Such analysis could improve coordination between sales, marketing, finance, and operations by showing when a seemingly attractive volume increase consumes scarce bottleneck capacity or raises expiry exposure. For Saudi manufacturers, this is particularly relevant during concentrated seasonal campaigns when many categories compete for the same lines, warehouses, transport slots, and packaging materials.

An eighth research need is stronger treatment of supplier uncertainty. Forecasting and scheduling models often focus on finished-goods demand while assuming ingredients and packaging arrive as planned. In practice, imported inputs may face port delays, geopolitical disruption, quality holds, or supplier capacity problems. A richer framework would combine demand scenarios with material-availability scenarios and estimate which constraints are most threatening to critical-SKU service. The resulting model could recommend alternative recipes where approved, earlier purchases, temporary capacity reallocation, or targeted safety stocks. This would make resilience planning more selective and economically defensible than simply increasing inventory across the portfolio.

A ninth direction concerns workforce and managerial adoption. Algorithms can propose schedules, but planners remain accountable for releasing them into environments where equipment conditions, labour skills, maintenance tasks, and customer commitments are imperfectly represented in data. Research should examine how planners interpret recommendations, which explanations increase confidence, and when human overrides improve or damage performance. Decision-support interfaces should expose constraint conflicts, uncertainty ranges, and the expected consequences of alternatives in language that cross-functional teams can use. Adoption studies should therefore measure trust calibration and decision quality rather than counting dashboard usage or model acceptance alone.

Finally, future Saudi case studies should compare integrated planning against credible operational baselines over long enough periods to include ordinary demand, promotional peaks, religious seasons, supply interruptions, and maintenance events. Evaluation should separate benefits caused by better forecasting from those caused by improved scheduling, while also measuring their interaction. A robust industrial study would report data preparation effort, model maintenance cost, computational time, planner overrides, service improvement, waste reduction, overtime, schedule stability, and recovery speed. This level of evidence would make it possible to judge whether the proposed architecture creates sustainable value beyond a successful analytics pilot and whether the same design principles transfer across dairy, beverage, bakery, frozen-food, meat-processing, and packaged-staple operations.

A tenth gap is the absence of standard operational benchmarks for integrated forecast-to-schedule systems. Future researchers should publish common definitions for demand, stockouts, freshness, schedule adherence, changeover loss, and service so that results can be compared across plants. Benchmark datasets should include event calendars, capacity constraints, inventory age, and disruption episodes rather than sales history alone. Shared evaluation protocols could then test whether an approach remains accurate, feasible, and stable when conditions change. For industry, such benchmarks would support more disciplined technology selection, reduce dependence on vendor demonstrations, and help management distinguish genuine capability from attractive visualization. They would also make replication easier across plants with different automation levels, product complexity, and planning maturity within the broader Saudi food manufacturing ecosystem today.

X. CONCLUSION

AI-assisted demand forecasting and production scheduling should be designed as one closed decision system for high-volume food manufacturing. Forecast models create value when they detect relevant demand signals, represent uncertainty, and remain calibrated under promotions, events, stockouts, and structural change. Scheduling models create value when they convert those signals into feasible sequences that respect capacity, sanitation, labour, materials, shelf life, freshness, and distribution requirements. The review shows that recent advances in machine learning, probabilistic forecasting, optimization, digital twins, and adaptive scheduling provide the technical components needed to connect these functions [1-24].

For Saudi Arabia, the strongest case for integration is strategic as well as operational. Better planning can improve availability, reduce waste, stabilize production, use domestic capacity more effectively, and shorten recovery after disruptions, supporting the broader objectives of food-processing localization and supply resilience [25-30]. The recommended pathway begins with trusted data and disciplined forecast governance, then adds finite-capacity scheduling, rolling replanning, scenario analysis, and network-level resilience. Success depends on transparent decision rules and cross-functional ownership rather than algorithmic sophistication alone. The proposed framework therefore positions AI as a planning capability embedded in manufacturing management, where predictive insight is continuously tested against physical constraints and actual execution outcomes.

REFERENCES

[1] F. Angizeh, H. Montero, A. Vedpathak, and M. Parvania, "Optimal production scheduling for smart manufacturers with application to food production planning," *Computers & Electrical Engineering*, vol. 84, Art. no. 106609, 2020, doi: 10.1016/j.compeleceng.2020.106609. ScienceDirect

[2] F. Kamhuber, T. Sobottka, B. Heinzl, J. Henjes, and W. Sihn, "An efficient hybrid multi-criteria optimization approach for rolling production smoothing of a European food manufacturer," *Computers & Industrial Engineering*, vol. 147, Art. no. 106620, 2020, doi: 10.1016/j.cie.2020.106620. ScienceDirect

[3] G. P. Georgiadis, B. Mariño Pampín, D. A. Cabo, and M. C. Georgiadis, "Optimal production scheduling of food process industries," *Computers & Chemical Engineering*, vol. 134, Art. no. 106682, 2020, doi: 10.1016/j.compchemeng.2019.106682. ScienceDirect

[4] C. G. Palacin, C. Vilas, A. A. Alonso, J. L. Pitarch, and C. de Prada, "Closed-loop Scheduling in a canned food factory," *IFAC-PapersOnLine*, vol. 53, pp. 10791–10796, 2020.

[5] S. Punia and S. Shankar, "Predictive analytics for demand forecasting: A deep learning-based decision support system," *Knowledge-Based Systems*, vol. 258, Art. no. 109956, 2022.

[6] D. Ni, Z. Xiao, and M. K. Lim, "A review and analysis of artificial intelligence methods for demand forecasting in supply chain management," *Procedia CIRP*, vol. 107, pp. 1126–1131, 2022.

[7] O. N. Arunkumar and D. Divya, "Deep learning techniques for demand forecasting: A review and future directions," *Information Resources Management Journal*, vol. 35, 2022.

[8] V. L. Migueis, A. Pereira, J. Pereira, and G. Figueira, "Reducing fresh fish waste while ensuring availability: Demand forecast using censored data and machine learning," *Journal of Cleaner Production*, vol. 359, Art. no. 131852, 2022.

[9] J. Huber, A. Gossmann, and H. Stuckenschmidt, "Daily retail demand forecasting using machine learning with emphasis on calendric special days," *International Journal of Forecasting*, vol. 36, pp. 1420–1438, 2020.

[10] K. Posch, C. Truden, P. Hungerländer, and J. Pilz, "A Bayesian approach for predicting food and beverage sales in staff canteens and restaurants," *International Journal of Forecasting*, vol. 38, pp. 321–338, 2022.

[11] F. Tangour, M. Nouiri, and R. Abbou, "Multi-Objective Production Scheduling of Perishable Products in Agri-Food Industry," *Applied Sciences*, vol. 11, Art. no. 6962, 2021.

[12] V. Solina and G. Mirabelli, "Integrated production-distribution scheduling with energy considerations for efficient food supply chains," *Procedia Computer Science*, vol. 180, pp. 797–806, 2021.

[13] M. E. Samouilidou, G. P. Georgiadis, and M. C. Georgiadis, "Food Production Scheduling: A Thorough Comparative Study between Optimization and Rule-Based Approaches," *Processes*, vol. 11, Art. no. 1950, 2023.

[14] R. Mousavi, M. Bashiri, and E. Nikzad, "Stochastic production routing problem for perishable products: Modeling and a solution algorithm," *Computers & Operations Research*, vol. 142, Art. no. 105725, 2022.

[15] V. Polotski, A. Gharbi, and J.-P. Kenne, "Production control in manufacturing systems with perishable products under periodic demand," *Journal of Manufacturing Systems*, vol. 63, pp. 288–303, 2022.

[16] A. Alvarez, P. Miranda, and S. U. K. Rohmer, "Production routing for perishable products," *Omega*, vol. 111, Art. no. 102667, 2022.

[17] Koulouris, N. Misailidis, and D. Petrides, "Applications of process and digital twin models for production simulation and scheduling in the manufacturing of food ingredients and products," *Food and Bioproducts Processing*, vol. 126, pp. 317–333, 2021.

[18] P. Maheshwari, S. Kamble, A. Belhadi, M. Venkatesh, and M. Z. Abedin, "Digital twin-driven real-time planning, monitoring, and controlling in food supply chains," *Technological Forecasting and Social Change*, vol. 195, Art. no. 122799, 2023.

[19] S. K. Jauhar, S. Harinath, V. Krishnaswamy, and S. K. Paul, "Explainable artificial intelligence to improve the resilience of perishable product supply chains by leveraging customer characteristics," *Annals of Operations Research*, vol. 354, pp. 103–142, 2025.

[20] N. Hübner, J. Caspers, V. C. Coroamă, *et al.*, "Machine-learning-based demand forecasting against food waste: Life cycle environmental impacts and benefits of a bakery case study," *Journal of Industrial Ecology*, vol. 28, pp. 1117–1131, 2024, doi: 10.1111/jiec.13528. Wiley

[21] K. Ait Ben Hamou, Z. Jarir, and S. Elfirdoussi, "Using machine learning for production scheduling problems in the supply chain: A review," *Computers & Industrial Engineering*, vol. 206, Art. no. 111243, 2025.

[22] A. Brintrup, E. Kosasih, P. Schaffer, G. Zheng, G. Demirel, and B. L. MacCarthy, "Digital supply chain surveillance using artificial intelligence: Definitions, opportunities and risks," *International Journal of Production Research*, vol. 62, pp. 4674–4695, 2024.

[23] H. Wu, J. Liu, and B. Liang, "AI-driven supply chain transformation in Industry 5.0: Enhancing resilience and sustainability," *Journal of the Knowledge Economy*, vol. 16, pp. 3826–3868, 2024.

[24] A. Rejeb, K. Rejeb, and A. Hassoun, "The impact of machine learning applications in agricultural supply chain: A topic modeling-based review," *Discover Food*, vol. 5, Art. no. 141, 2025.

[25] Saudi Vision 2030, *Annual Report 2025*. Riyadh, Saudi Arabia, 2025.

[26] National Industrial Development and Logistics Program, *Delivery Plan: Food Processing and Industrial Development Priorities*. Riyadh, Saudi Arabia, 2024.

[27] Saudi Vision 2030, *National Industrial Development and Logistics Program: Food Processing Sub-Sector Strategy*. Riyadh, Saudi Arabia, 2022.

[28] General Authority for Statistics, *Food Security Statistics 2024*. Riyadh, Saudi Arabia, 2024.

[29] Food and Agriculture Organization of the United Nations, *GIEWS Country Brief: Kingdom of Saudi Arabia*. Rome, Italy, 2024.

[30] F. Avishan, İ. Yanıkoğlu, and M. Soysal, "Adaptive optimization approach for production and distribution planning of perishable food products under demand uncertainty," *Annals of Operations Research*, 2025, doi: 10.1007/s10479-025-06552-5. Springer