THE COMPUTERTECH
(An International Peer Review Journal)
VOLUME 12; ISSUE 1 (JAN-JUNE); (2026)
WEBSITE: THE COMPUTERTECH

Readiness, Adoption and Governance of AI-Driven Financial Planning in Small and Medium-Sized Enterprises: A Staged Implementation Framework

[1] Mariam Borg, [2] Maya Hakobyan
[1] University of Ljubljana, SLOVENIA
[2] Technical University of Denmark, DENMARK

Abstract
Artificial intelligence is beginning to bring capabilities once reserved for large corporate finance teams, such as machine-learning cash-flow forecasting, scenario analysis and risk-aware decision support, within reach of small and medium-sized enterprises (SMEs). Recent work has set out an AI-driven financial planning and analysis (FP&A) approach for technology-enabled SMEs that combines these elements [1]. Whether such systems deliver value, however, depends less on algorithms than on whether a firm is ready to adopt them and able to govern them once in use. This article develops a staged implementation framework for AI-driven FP&A in SMEs. Drawing on technology-adoption theory, forecasting research and emerging AI risk-management guidance, it proposes a reference architecture, a readiness assessment across five dimensions, a five-level maturity model, an eighteen-month phased roadmap with explicit decision gates, a catalogue of risks and controls, and a set of adoption metrics. The framework emphasizes data foundations before modelling, shadow-mode testing against simple benchmarks, human review of every material decision, and transparent communication of forecast uncertainty. We test the data and baseline gates on real weekly revenue of 327 small online sellers, implementing the machine-learning forecasting engine of the referenced FP&A approach. Sellers scored as more data-ready were more likely to pass the baseline gate (45 percent against 37 percent for the least ready) and received about half as many false shortfall alerts, but in every readiness tier most sellers would have been better served by a simple benchmark, which is the outcome the gate is designed to catch. The framework is intended to help SME owners, advisers, software providers and lenders move from interest in AI-driven planning to its safe and productive use.

Keywords: AI-driven FP&A; SMEs; technology adoption; cash-flow forecasting; digital readiness; maturity model; AI governance; decision support

Introduction

Small and medium-sized enterprises account for most firms and a large share of employment, yet they typically manage their finances with limited staff, spreadsheets and backward-looking accounts. Transaction data on hundreds of thousands of US small businesses show that a typical firm holds only a few weeks of cash buffer [2], so the ability to look ahead matters. International policy work has identified digital adoption as a central lever for SME productivity, while noting that smaller firms lag larger ones in adopting data-intensive technologies [3].

AI-driven financial planning and analysis offers a way to close part of this gap. The approach described in [1] combines machine-learning cash-flow forecasting with risk-aware decision support for technology-enabled operations. Such systems can automate data collection, produce rolling forecasts, run scenarios and alert owners to emerging risks. Yet the experience of data-driven decision-making in firms shows that technology alone rarely changes outcomes; the complementary investments in data, skills and processes determine whether adoption pays off [4].

This article addresses the practical question that follows: how should an SME prepare for, adopt and govern an AI-driven FP&A system? We contribute a reference architecture, a readiness assessment, a maturity model, a phased roadmap with decision gates, a catalogue of risks and controls, and adoption metrics, and we test the data-readiness assessment and the baseline gate on real revenue data from 327 small sellers. The framework is intended for owners, external accountants and advisers, software providers, and lenders who increasingly rely on SME financial data.

Background and Related Work

Technology adoption has been studied extensively at both the individual and the firm level. At the individual level, perceived usefulness and ease of use shape whether people adopt a system [5], and the unified theory of acceptance adds social influence and facilitating conditions [6]. At the firm level, the technology-organization-environment (TOE) framework holds that adoption depends on the characteristics of the technology, the resources and structure of the organization, and the external environment of customers, suppliers and regulators [7][8]. These perspectives are directly relevant to SMEs, where the owner is often both decision-maker and primary user.

Forecasting research provides two further cautions. First, sophisticated machine-learning methods do not automatically outperform simple statistical benchmarks, so new tools must be tested against them [9][10]. Second, forecasts in organizations are routinely adjusted by human judgment, which can improve or damage accuracy [11]. Users also tend to abandon algorithms after seeing them err [12], although allowing modest adjustments sustains use [13]. An adoption framework must therefore treat human review as a design feature rather than an afterthought.

Finally, machine-learning systems carry ongoing maintenance costs arising from data dependencies and changing conditions [14], and data quality has many dimensions beyond accuracy, including timeliness, completeness and relevance [15]. The NIST AI Risk Management Framework organizes responsible AI practice around four functions, govern, map, measure and manage, which scale down usefully to small organizations [16]. Table 1 contrasts traditional SME financial planning with the AI-driven approach of [1].

Table 1. Traditional and AI-driven financial planning in SMEs.

Capability	**Traditional SME practice**	**AI-driven FP&A**
Data collection	Manual exports and spreadsheet entry	Automated feeds from bank, ledger, invoicing and payroll systems
Forecast cadence	Annual budget, occasional updates	Rolling weekly or monthly forecasts
Forecast form	Single-number projection	Median with uncertainty range

Capability	Traditional SME practice	AI-driven FP&A
Scenario analysis	Ad hoc, limited	Systematic what-if analysis on demand
Risk signals	Noticed after the fact	Early alerts on cash shortfalls and covenant pressure
Decision support	Owner intuition	Recommendations reviewed by owner or finance lead

Reference Architecture

Figure 1 presents a reference architecture for AI-driven FP&A in an SME. Data from bank feeds, the accounting ledger, invoicing and payables, and payroll or point-of-sale systems pass through a data layer that cleanses, maps and checks them. A forecasting engine produces cash-flow and revenue forecasts, which feed scenario and risk analysis. Decision support turns these into alerts, plans and budgets, which are reviewed by the owner or finance lead before action. Overrides and actual results feed back into data and model monitoring.

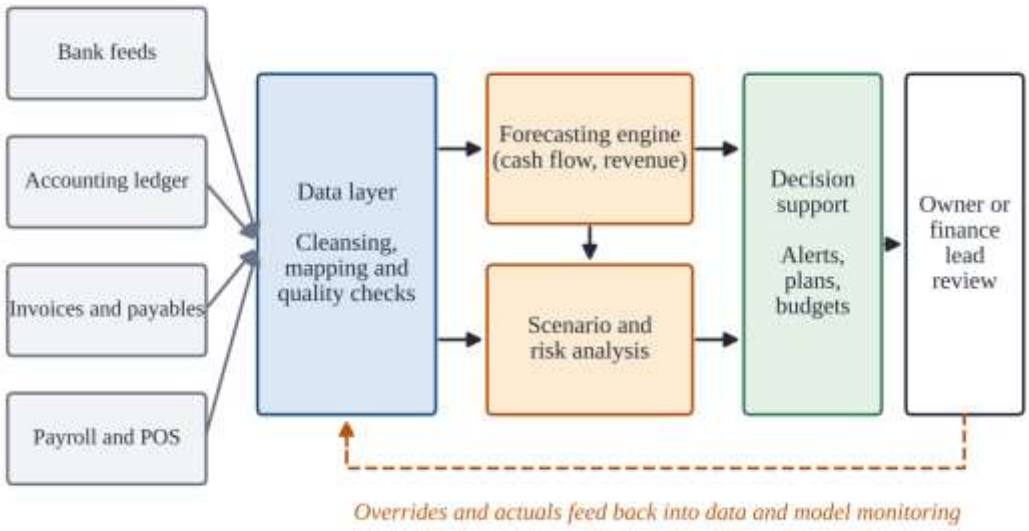


Figure 1. Reference architecture for AI-driven FP&A in an SME [authors' design].

Two design principles follow from the literature. First, the data layer deserves as much attention as the model, because most failures in small-firm analytics arise from incomplete or inconsistent data [15][14]. Second, the final box, human review, is not optional: it is where the owner's knowledge of contracts and customers corrects the model, and where accountability for decisions remains [11][16].

Readiness Assessment

Before investing, an SME should assess its readiness across five dimensions that extend the TOE framework [7] with two dimensions particularly important for analytics: data and people. Table 2 lists indicators for each. A firm need not score highly on every dimension before starting, but a low score on data readiness is a strong signal to begin with Phase 1 of the roadmap rather than with modelling.

Table 2. Readiness dimensions and indicators for AI-driven FP&A.

Dimension	**Indicators of readiness**	**Warning signs**
Technology	Cloud accounting; bank feeds; integration with invoicing and payroll	Desktop-only software; manual bank reconciliation
Data	At least two years of clean, categorized transactions; consistent chart of accounts	Frequent miscoding; missing periods; personal and business funds mixed
Organization	Named finance owner; monthly close; willingness to change routines	No regular close; decisions made without reference to accounts
People	Owner or staff able to interpret ranges and scenarios	Discomfort with uncertainty; reliance on a single external accountant
Environment	Lender, investor or customer demand for forecasts; supportive adviser	No external need for planning; regulatory constraints on data sharing

Applying the Assessment: An Illustrative Example

To show how the assessment guides decisions, consider a hypothetical wholesale distributor with 25 employees, cloud accounting software and a single part-time bookkeeper. The example is illustrative and does not describe a real firm. Each dimension in Table 2 is scored from 1 (not ready) to 5 (fully ready) by the owner and an external adviser together. Table 3 records the result.

Table 3. Illustrative readiness scoring for a hypothetical 25-employee distributor.

Dimension	**Score (1–5)**	**Evidence**	**Implication**
Technology	4	Cloud ledger with bank feeds; invoicing integrated	No major barrier
Data	2	Inconsistent coding of supplier payments; one missing quarter	Start at Phase 1; clean two years of data
Organization	3	Monthly close, but often late	Set a ten-day close target
People	2	Owner unfamiliar with forecast ranges	Adviser-led training before Phase 4

Dimension	Score (1–5)	Evidence	Implication
Environment	4	Bank requests quarterly forecasts for credit line renewal	Strong incentive to adopt

Although the firm's technology is adequate and its bank is pressing for forecasts, its low data and people scores mean that moving straight to machine learning would be premature. The assessment points instead to a data-cleaning project and a baseline rolling forecast, with machine-learning forecasting introduced only after the data gate is passed. In this way the assessment protects the firm from investing in a model before the foundations exist to make it useful [4].

Maturity Model

Adoption is a progression rather than a switch. Figure 2 sets out five maturity levels. At Level 1, forecasts are built manually in spreadsheets. Level 2 connects data feeds and introduces rolling forecasts. Level 3 adds a machine-learning baseline forecast reviewed by a person. Level 4 makes forecasts probabilistic and adds scenarios and alerts, the risk-aware decision support emphasized in [1]. Level 5 brings the system under formal governance, with monitoring, documentation and, where useful, integration with lenders.

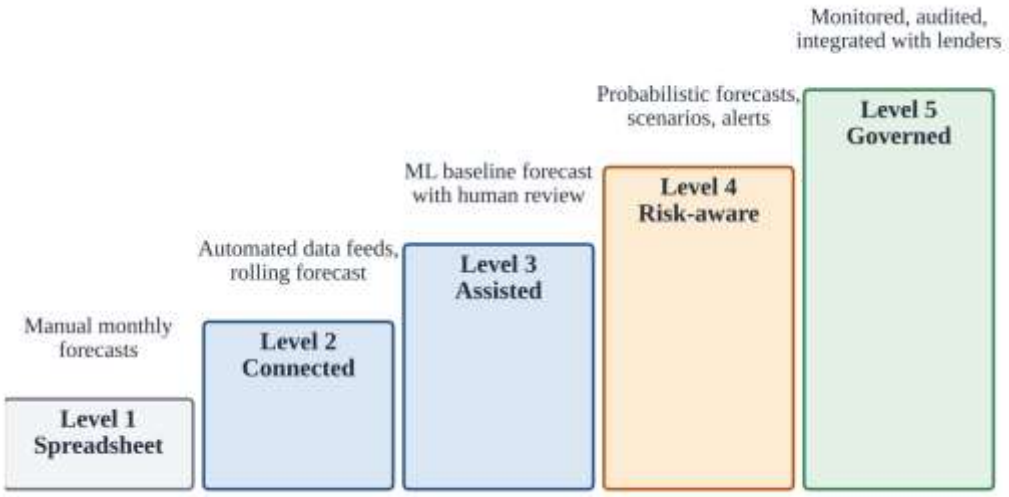


Figure 2. Five-level maturity model for AI-driven FP&A in SMEs [authors' design].

Table 4 states the criteria for moving from one level to the next. The criteria are deliberately concrete, so that an owner or adviser can judge progress without specialist knowledge.

Table 4. Criteria for progressing between maturity levels.

Transition	Criteria to progress
Level 1 to 2	Bank, ledger and invoicing data flow automatically; monthly close completed within ten working days

Transition	Criteria to progress
Level 2 to 3	Rolling forecast maintained for six months; baseline accuracy of a simple seasonal method recorded [10]
Level 3 to 4	Machine-learning forecast beats or matches the baseline in shadow mode; owner reviews each forecast [9]
Level 4 to 5	Uncertainty ranges shown and checked against outcomes; alerts tested; documented review and override process [16]

Phased Implementation Roadmap

Figure 3 sets out an indicative eighteen-month roadmap. Phase 1 builds the data foundation. Phase 2 introduces a baseline rolling forecast using simple methods, which also creates the benchmark against which machine learning will be judged. Phase 3 runs a machine-learning forecast in shadow mode, producing forecasts that are recorded but not acted on. Phase 4 introduces assisted use, in which the owner sees and may adjust the machine-learning forecast [13]. Phase 5 scales the system and brings it under governance.

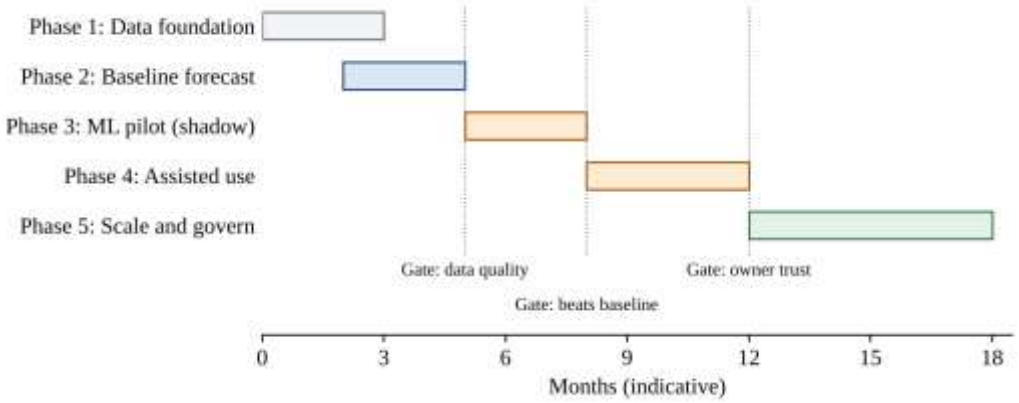


Figure 3. Indicative phased roadmap with decision gates [authors' design].

Each phase ends at a gate. The data-quality gate requires that feeds are complete and reconciled. The baseline gate requires that the machine-learning forecast at least matches the simple benchmark over the shadow period; if it does not, the firm should continue with the benchmark rather than adopt a more complex tool for its own sake [9]. The trust gate requires that the owner understands the forecast ranges and has used them in at least one real decision.

Use Cases

AI-driven FP&A supports a range of decisions. Table 5 lists common use cases, the data each requires and the value it offers. Cash-flow forecasting is usually the best starting point because the data are readily available from bank feeds and the value is immediate [2].

Table 5. Use cases for AI-driven FP&A in SMEs.

Use case	Data required	Value to the firm
13-week cash forecast	Bank transactions; receivables and payables	Early warning of shortfalls; timing of payments
Revenue forecasting	Sales history; seasonality; pipeline	Staffing and inventory planning
Scenario analysis	Forecast model plus assumptions	Testing price changes, new hires or lost customers
Working-capital optimization	Invoices, payment terms and history	Faster collection; better supplier terms
Lender reporting	Forecasts and actuals	Stronger credit applications; earlier conversations with lenders

Risks and Controls

AI-driven planning introduces risks that small firms may not anticipate. Table 6 pairs the principal risks with proportionate controls. The controls are framed for small organizations and map to the govern, map, measure and manage functions of the NIST framework [16].

Table 6. Principal risks of AI-driven FP&A in SMEs and proportionate controls.

Risk	Example	Proportionate control
Data errors	Miscoded transactions distort forecasts [15]	Monthly reconciliation; automated data-quality flags
Overconfidence	Narrow ranges understate the chance of a shortfall	Show uncertainty ranges and check them against outcomes
Automation bias	Owner accepts forecasts without scrutiny	Require review and a recorded reason for major decisions
Algorithm aversion	Owner abandons tool after one bad forecast [12]	Allow bounded adjustments; explain errors within stated ranges
Model decay	Forecast quality falls as the business changes [14]	Quarterly accuracy and coverage review; retrain when needed
Vendor dependence	Provider changes or withdraws the service	Export rights for data and forecasts; documented fallback method
Privacy and security	Sensitive financial data shared with third parties	Least-privilege access; review of provider security terms

Measuring Adoption Success

Success should be measured by outcomes rather than by installation. Table 7 proposes metrics in three groups: forecast quality, use, and business outcomes. Forecast quality should include both accuracy relative to a simple benchmark and the reliability of uncertainty ranges, since an accurate but overconfident forecast can mislead [17][18].

Table 7. Adoption metrics for AI-driven FP&A.

Group	Metric	Example target
Forecast quality	Error relative to seasonal-naive benchmark [17]	At or below 1.0
Forecast quality	Share of outcomes within stated 80% range [18]	Between 75% and 85%
Use	Share of months with forecast reviewed and signed off	100%
Use	Share of forecasts adjusted, with reasons recorded	Tracked; reasons reviewed quarterly
Business outcome	Number of unplanned overdrafts or late payments	Falling over time
Business outcome	Lead time of shortfall alerts before the event	At least four weeks

Testing the Data and Baseline Gates on Real Seller Data

Two parts of the framework make claims that data can test: that a low data-readiness score signals a firm should not yet rely on machine-learning forecasts, and that the baseline gate, which requires a model to at least match a simple benchmark during the shadow period, protects firms from premature adoption. We tested both on the public Olist dataset of orders on a Brazilian online marketplace[19], which records every sale by thousands of small sellers. Weekly revenue was computed for 327 sellers that traded from January 2017 and had sales in at least 15 of their first 52 weeks. Median weekly revenue was BRL 134. Revenue is a proxy for cash inflow only; outflows are not observed.

The forecasting engine followed the approach in [1]: a global gradient-boosted quantile model (LightGBM[20]) produced a median and an 80 percent interval for revenue over the next four weeks, and the interval was conformally calibrated[21]. It was compared with a naive forecast (the mean of the last four weeks) and a seasonal-naive forecast (the same weeks a year earlier). Models were trained on forecast origins in weeks 20 to 51 of 2017, calibrated on weeks 52 to 61, and evaluated on origins in weeks 62 to 80 (the shadow period). Data readiness was scored from information available before the shadow period, using three indicators from the data dimension of Table 2: the share of weeks with no recorded sales (missing periods), the coefficient of variation of weekly revenue (consistency), and the number of active weeks (history length). Sellers were placed in low, medium or high tiers by their combined tertile scores.

Table 8. Data-readiness tiers and baseline-gate outcomes over the shadow period (327 Olist sellers; MAE relative to the naive forecast).

Readiness tier	Sellers	Median share of zero-sale weeks	Median CV	ML relative MAE	Seasonal-naive relative MAE	Sellers passing baseline gate	Conformal 80% coverage	False shortfall alerts
Low	106	62%	1.71	0.90	1.86	37%	84%	33%
Medium	83	44%	1.24	1.07	2.14	37%	81%	30%
High	138	12%	0.87	0.97	1.88	45%	82%	16%

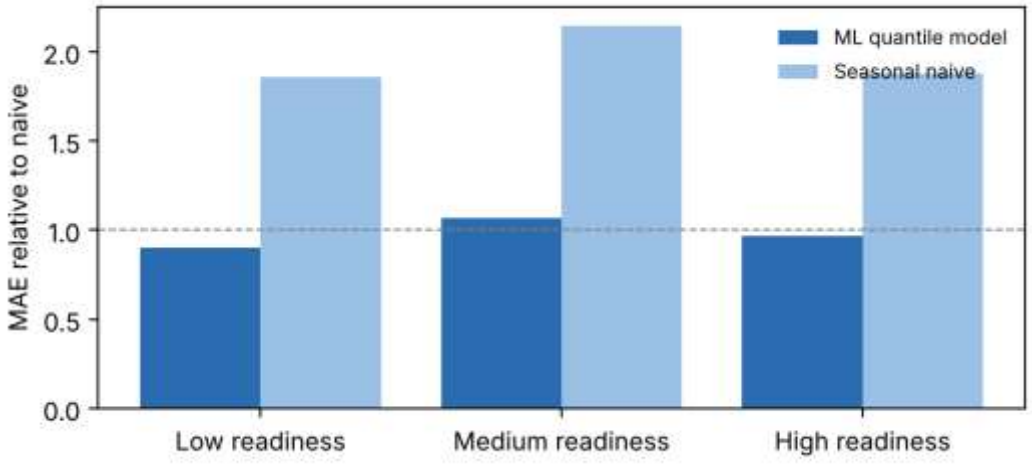


Figure 4. Forecast error relative to the naive method by data-readiness tier, machine-learning model versus seasonal naive [authors' analysis].

The results support the gates, though not in the simplest way. Averaged over each tier, the machine-learning model's error relative to the naive forecast did not improve with readiness (0.90, 1.07 and 0.97 for the low, medium and high tiers), because the naive benchmark itself performs worst on the sparse, erratic series of low-readiness sellers. At the level of the individual firm, which is the level at which the gate is applied, readiness did matter. The share of sellers whose machine-learning forecasts beat the better of the two simple benchmarks over the shadow period rose from 37 percent in the low tier to 45 percent in the high tier. Shortfall alerts based on the median forecast raised false alarms in 33 percent of safe cases for low-readiness sellers but 16 percent for high-readiness sellers, which matters for the trust gate, since owners learn to ignore frequent false alarms. Conformal calibration delivered close to nominal interval coverage in every tier (81 to 84 percent). The most important result is that in every tier a majority of sellers failed the baseline gate. Even among the most data-ready sellers, 55 percent would have been better served over the shadow period by a naive or seasonal forecast, and the seasonal-naive method was poor for all tiers (relative MAE around 2.0) because most sellers were growing. Without the gate, a typical small seller would have adopted a machine-learning tool that did not outperform a four-week average. The finding is specific to young, fast-growing online sellers with revenue-only data. Firms with longer histories

and full ledger data may pass the gate more often, which is exactly why the framework requires each firm to run the test on its own data rather than relying on a vendor's average.

Discussion

The framework rests on a simple observation: the value of AI-driven FP&A lies in better decisions, and better decisions require reliable data, honest uncertainty and human judgment. The approach in [1] supplies the analytical engine. The readiness assessment and maturity model help an SME judge where it stands; the roadmap and gates prevent premature reliance on unproven models; and the risk controls and metrics keep the system accountable once in use. The test on real seller data shows that the baseline gate is not a formality: most small sellers in the sample would have failed it, and data readiness raised the pass rate without making the gate unnecessary.

Advisers and lenders have important roles. External accountants can carry out readiness assessments and supervise early phases, compensating for limited in-house skills. Lenders that receive forecast data can encourage adoption by valuing well-governed forecasts, while treating unvalidated forecasts with appropriate caution. Software providers can lower adoption barriers by building benchmark comparisons, uncertainty ranges and override logs into their products by default.

Table 9. Roles of stakeholders in SME adoption of AI-driven FP&A.

Stakeholder	Main role	Contribution to safe adoption
Owner or finance lead	Reviews forecasts and makes decisions	Supplies context; records overrides and reasons
External accountant or adviser	Assesses readiness; supervises early phases	Bridges skills gap; checks data quality
Software provider	Supplies data integration and models	Builds in benchmarks, uncertainty ranges and override logs
Lender	Uses forecasts in credit relationships	Rewards well-governed forecasts; treats unvalidated ones with caution

Table 9 summarizes these roles. Their combined effect is to spread the work of governance so that no single party, least of all a time-pressed owner, carries it alone.

Conclusion

AI-driven financial planning can give SMEs forward-looking insight that was once available only to much larger firms. Realizing that promise requires more than a capable model. This article has proposed a staged framework, comprising a reference architecture, readiness assessment, maturity model, gated roadmap, risk controls and adoption metrics, that helps SMEs build the data foundations, test models against simple benchmarks, keep people in control of decisions and govern their systems over time. Tested on 327 small online sellers, the data-readiness score predicted which firms were more likely to benefit from machine-learning forecasts and to receive fewer false

alerts, and the baseline gate would have stopped most of them from adopting a model that did not beat a simple benchmark. Combined with the analytical approach of [1], the framework offers a practical path from interest in AI-driven planning to its safe and productive use.

References

[1] Karim, N. B., & Seitz, J. (2023). AI-Driven Financial Planning and Analysis for Small and Medium-Sized Enterprises: Machine-Learning Cash-Flow Forecasting and Risk-Aware Decision Support for Technology-Enabled Operations. *International Journal of Modern Computing*, 6(1), 158-176.

[2] Farrell, D., & Wheat, C. (2016). *Cash is king: Flows, balances, and buffer days—Evidence from 600,000 small businesses*. JPMorgan Chase & Co. Institute.

[3] OECD. (2021). *The digital transformation of SMEs*. OECD Publishing.

[4] Brynjolfsson, E., & McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, *106*(5), 133–139.

[5] Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, *13*(3), 319–340.

[6] Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, *27*(3), 425–478.

[7] Tornatzky, L. G., & Fleischer, M. (1990). *The processes of technological innovation*. Lexington Books.

[8] Oliveira, T., & Martins, M. F. (2011). Literature review of information technology adoption models at firm level. *Electronic Journal of Information Systems Evaluation*, *14*(1), 110–121.

[9] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLoS ONE*, *13*(3), e0194889.

[10] Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts.

[11] Fildes, R., Goodwin, P., Lawrence, M., & Nikolopoulos, K. (2009). Effective forecasting and judgmental adjustments: An empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting*, *25*(1), 3–23.

[12] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, *144*(1), 114–126.

[13] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, *64*(3), 1155–1170.

[14] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. In *Advances in Neural Information Processing Systems* (Vol. 28, pp. 2503–2511).

[15] Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems, 12*(4), 5–33.

[16] National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). U.S. Department of Commerce.

[17] Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting, 22*(4), 679–688.

[18] Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B, 69*(2), 243–268.

[19] Olist. (2018). *Brazilian e-commerce public dataset by Olist* [Data set]. Olist. https://github.com/olist/work-at-olist-data

[20] Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems 30* (pp. 3146-3154).

[21] Romano, Y., Patterson, E., & Candès, E. (2019). Conformalized quantile regression. In *Advances in Neural Information Processing Systems 32* (pp. 3543-3553).