

International Journal of Scientific Interdisciplinary Research

Volume: 1; **Issue:** 1
Pages: 316–334

Predictive Cash Flow Forecasting Using Deep Learning and ERP Transaction Data in Mid-Market Manufacturing Firms

Samia Hossain Swarnali[1]; H M Mahir Uddin[2]

[1]. *Senior Financial Analyst, GETCO Business Solutions Limited, Dhaka, Bangladesh ;*
Email: shossain.samia@gmail.com
[2]. *Relationship Officer, Transaction Banking Department, BRAC Bank Limited, Dhaka, Bangladesh*
Email: mahir.uddin@pace.edu

Doi: 10.63125/mdsdab78
Received: *25 April 2022;* ***Revised:*** *29 May 2022;* ***Accepted:*** *13 June 2022;* ***Published:*** *30 June 2022*

Abstract

Mid-market manufacturing firms occupy an awkward middle ground in corporate finance. They are large enough to generate transaction volumes that exceed the capacity of spreadsheet-driven treasury practices, yet too small to maintain the dedicated quantitative finance teams that large enterprises deploy for liquidity management. The result is a persistent gap between the cash-flow visibility these firms need and the forecasting tools they actually use. This paper develops a conceptual and methodological framework for closing that gap by combining deep learning models with the granular transaction data already captured inside enterprise resource planning (ERP) systems. We argue that the ERP layer — accounts receivable, accounts payable, sales orders, production schedules, inventory movements, and the general ledger — constitutes an underexploited, high-frequency data source that is naturally suited to sequence-modeling architectures such as long short-term memory (LSTM) networks and attention-based transformers. Drawing on a structured review of more than sixty sources spanning financial forecasting, machine learning, working-capital theory, ERP analytics, and SME technology adoption, we propose a forecasting pipeline that ingests ERP transaction streams, engineers liquidity-relevant features, trains a hybrid recurrent–attention model, and surfaces explainable, decision-ready forecasts through a treasury dashboard. Because no proprietary firm dataset was available for this study, the empirical section presents a clearly labeled illustrative simulation rather than results from real operating data. In that simulation, deep learning architectures reduce forecasting error relative to a seasonal ARIMA baseline by margins consistent with the comparative literature. We discuss why interpretability, data governance, and modest implementation cost are decisive for adoption in resource-constrained firms, and we close with recommendations and an explicit account of the study's limitations. The contribution is a deployable blueprint and research agenda rather than a validated empirical claim.

Keywords

Cash Flow Forecasting; Deep Learning; Enterprise Resource Planning (ERP); Mid-Market Manufacturing; Explainable Artificial Intelligence.

INTRODUCTION

Cash is the constraint that binds every other decision a manufacturing firm makes. A factory can hold a healthy order book, a defensible gross margin, and a growing customer base, and still fail because the timing of cash inflows and outflows drifts out of alignment (Hiebl et al., 2017). Cash flow forecasting, the modeling of an entity's future liquidity over a defined horizon, exists precisely to manage that timing risk. In the short term it governs working-capital management; over longer horizons it informs asset and liability management. For large corporations, forecasting is supported by treasury departments, dedicated analysts, and quantitatively sophisticated planning systems. For mid-market manufacturers, firms with revenues roughly between fifty million and one billion dollars, the same task is frequently performed in spreadsheets maintained by a handful of finance staff, using historical averages and static assumptions that struggle to keep pace with fast-changing market and operational conditions. This is not a marginal segment (Jagoda & Samaranayake, 2017). Mid-market manufacturers form the backbone of industrial supply chains, and their financial fragility is well documented. Working-capital management is repeatedly shown to be a central driver of performance and survival for small and medium-sized manufacturers whose current assets account for a large share of total investment. Short-term credit constraints turn out to be at least as important as long-term ones in shaping these firms' investment decisions, because cash credit constraints force managers to forgo opportunities in order to finance working-capital needs. A longer cash conversion cycle is consistently associated with lower profitability in manufacturing SMEs, and the same cycle reduces the likelihood and volume of exports for financially constrained firms. In short, the firms least equipped to forecast cash flow accurately are precisely the firms for which accurate forecasting matters most (Mayeh et al., 2016).

At the same time, two developments have quietly converged to make a different approach feasible. The first is the maturation of enterprise resource planning systems among mid-market manufacturers. Modern ERP platforms now capture, in a single integrated repository, the financial and operational events that ultimately determine cash position: customer invoices and their payment behavior, supplier bills and payment terms, sales orders, production schedules, inventory movements, and the general ledger. Industry surveys report that analytics platforms and ERP systems are among the software investments manufacturers say deliver the quickest impact, and that a large majority of manufacturing leaders are pursuing a data-first strategy (Jagoda & Samaranayake, 2017). Business intelligence and analytics are increasingly treated as essential, not optional, features of SME manufacturing ERP. The transaction-level exhaust of these systems is, in effect, a high-frequency time series of the firm's liquidity—yet it is rarely modeled as one. The second development is the demonstrated superiority of deep learning for nonlinear, non-stationary financial time series. Recurrent neural networks with memory are specifically designed to learn dependencies across long sequences (Jagoda & Samaranayake, 2017; Lin et al., 2005). A substantial body of comparative work reports that LSTM models outperform classical autoregressive integrated moving average (ARIMA) models for financial prediction, chiefly because they capture nonlinear relationships that linear models cannot. Deep learning with LSTM networks has been applied successfully to financial market prediction at scale, and a systematic review of the 2005–2019 literature documents the breadth of deep learning applications across financial time series tasks. More recently, attention-based transformer architectures—originating in natural language processing—have been adapted to forecasting, offering global attention over long input windows.

The intersection of these two developments defines the opportunity this paper addresses (Mayeh et al., 2016; Mayeh et al., 2014). The central proposition is straightforward: mid-market manufacturers can materially improve cash-flow forecast accuracy by treating ERP transaction streams as the primary input to a deep learning forecasting pipeline, rather than aggregating that data into coarse monthly figures and forecasting with spreadsheets. Early industry work already points in this direction, reporting that AI-driven forecasting integrated with ERP systems and real-time transaction data

improves prediction granularity and responsiveness. Academic models combining neural networks with LSTM layers for corporate cash flow prediction have likewise been proposed (Morris & Venkatesh, 2010; Nah et al., 2004).

Yet a proposition is not a solution. Three obstacles stand between the opportunity and its realization in the mid-market. First is adoption capacity: SMEs face well-documented barriers to AI implementation—limited budgets, scarce technical skills, legacy IT, and data silos—that large enterprises do not. Second is interpretability: a treasurer will not act on a forecast they cannot interrogate, and the opacity of deep models is a recognized barrier to their acceptance in finance. Third is governance: deploying AI in financial forecasting raises real obligations around data privacy, auditability, and bias. A credible framework must therefore address not only model accuracy but also the organizational and ethical conditions under which a mid-market firm could actually use it. This paper makes four contributions. It synthesizes the dispersed literatures on financial forecasting, ERP analytics, and working-capital theory into a single problem statement for the mid-market segment. It specifies an end-to-end forecasting pipeline—from ERP feature engineering through a hybrid LSTM–attention model to an explainable treasury dashboard. It presents an illustrative simulation that demonstrates the expected shape of results and the diagnostic logic a real deployment would follow, while being explicit that the figures are synthetic. And it translates the analysis into governance-aware recommendations calibrated to the resource realities of mid-market manufacturers. The remainder of the paper proceeds through a literature review, a methodology, illustrative findings, a discussion, conclusions, recommendations, and a candid statement of limitations.

Theoretical Framing

Three theoretical lenses motivate the framework and recur throughout the paper. The first is the pecking-order theory of financing, in which firms prefer internal funds to external finance because of information asymmetry between managers and outside investors. Under this logic, the ability to forecast and manage internal cash is not a back-office convenience but a determinant of a firm's strategic latitude: a firm that can reliably anticipate its own liquidity needs less costly external capital and can deploy internal funds more confidently. For mid-market manufacturers, whose access to external finance is often constrained, the marginal value of an accurate internal forecast is correspondingly high. The second lens is the cash-conversion-cycle view of working-capital management, which links the timing of receivables, payables, and inventory to profitability. Because the cycle is a timing construct, it is intrinsically a forecasting problem: managing it well requires anticipating when cash will arrive and depart, not merely recording when it did. This is the conceptual bridge between the finance literature and the time-series-modeling literature—the cash conversion cycle is, in effect, a quantity that a sequence model is built to predict (Pan & Jang, 2016; Park et al., 2007). The third lens is the technology–organization–environment framework, often combined with diffusion-of-innovations theory, which explains why otherwise capable technologies are or are not adopted. This lens insists that the technological merit of a forecasting pipeline is only one of three determinants of its success; organizational readiness (skills, culture, leadership) and environmental factors (competitive pressure, vendor ecosystem, regulation) matter just as much. Holding these three lenses together—financing theory, working-capital theory, and adoption theory—is what distinguishes a deployable framework from a purely technical proposal.

The Mid-Market Manufacturing Context

Three features of the segment make it the natural setting for this framework. First, manufacturing cash flow is structurally complex: it couples a procure-to-pay cycle (raw materials, work-in-progress, supplier terms) with an order-to-cash cycle (sales orders, production lead times, customer payment behavior), so the timing of cash is the product of several interacting operational processes rather than a simple revenue stream. This is exactly the kind of multi-driver, nonlinear timing problem at which sequence models excel. Second, the segment sits at a capability threshold. Mid-market firms are large enough to run a full ERP system that captures the necessary transaction data, yet small enough that

they rarely employ the data-science staff that would make a bespoke forecasting model routine. They therefore have the raw material but not the modeling layer—the precise gap a managed, explainable pipeline can fill. Third, the financial stakes are high and asymmetric: working-capital efficiency strongly conditions profitability and survival in this segment, and short-term credit constraints have outsized effects on investment, so the marginal return to better forecasting is unusually large.

Figure 1: The cash conversion cycle (CCC = DIO + DSO − DPO).



LITERATURE REVIEW

The literature relevant to this study spans four largely separate conversations: the theory and practice of cash flow forecasting; the comparative performance of statistical and deep learning models for financial time series; the role of ERP systems as analytical infrastructure; and the financial behavior and technology-adoption constraints of mid-market and SME manufacturers. Each conversation is mature in isolation, but their intersection—deep learning applied to ERP-sourced cash-flow data in resource-constrained firms—remains thinly developed. This section reviews each strand and identifies the gap the present study addresses.

Cash Flow Forecasting: Foundations and Limits of Traditional Practice

Cash flow forecasting models a firm's future liquidity, with short-horizon forecasts oriented toward working-capital management and longer-horizon forecasts toward asset and liability management (Dechow et al., 1998). Traditional corporate practice relies on a mix of direct methods, projecting individual receipts and disbursements, and indirect methods derived from pro forma financial statements. The accounting literature has long studied the quality of such forecasts: it has been examined analysts' cash flow forecasts and found them materially less accurate and more naive than their earnings forecasts, underscoring how difficult cash-flow prediction is even for sophisticated professionals. Discounted cash flow valuation, the dominant framework linking forecasts to value, depends critically on the quality of the underlying projections. The structural weakness of conventional practice is its reliance on fixed assumptions, historical averages, and spreadsheets that fail to keep pace with rapidly changing conditions and that integrate poorly with operational systems. Even recent machine-learning-enhanced valuation work notes that cash flow forecasts traditionally rely on rigid assumptions, and proposes integrating LSTM with Monte Carlo simulation to bypass those limitations. The emerging consensus is that forecasting accuracy is the binding constraint on a chain of downstream decisions—investment strategy, dividend policy, and debt-covenant compliance —and that traditional statistical models struggle to capture the nonlinear dynamics those decisions depend on.

From ARIMA to Deep Learning: Comparative Evidence

The classical workhorse for time-series forecasting is the ARIMA family formalized by Box and Jenkins, valued for mathematical simplicity and for performing well with relatively small datasets. Its limitation is structural: being linear, it captures trends and seasonality but not the nonlinear patterns common in financial data. The newest frontier is the transformer. Self-attention enables a model to weigh all time steps simultaneously rather than sequentially, which improves both parallelism and the modeling of long-range dependencies (Golam & Amir, 2022; Abdur & Iftekhar, 2021). Surveys document rapid adoption of transformer variants for time series, and dual-attention transformer architectures have been proposed specifically for financial prediction. A sobering counter-current, however, warns against uncritical adoption: Ughi, Lomurno, and Matteucci show that simplifying transformer forecasters often improves performance, and that shallow models without attention can match the state of the art—arguing that a simple baseline must always be used to verify any claimed gain. This caution is directly relevant to mid-market firms, where simpler models are also cheaper to deploy and maintain (Atif & Murad, 2022; Binayan & Shakhawat, 2022). It is worth dwelling on why the linear–nonlinear distinction matters so concretely for cash flow. ARIMA models decompose a series into autoregressive and moving-average components around a stationary mean, which works well when the data-generating process is approximately linear and the relationships are stable. Manufacturing cash flow violates these assumptions in characteristic ways: customer payment behavior is bursty and threshold-driven; production schedules introduce step changes in payables; and demand shocks create regime changes that a fixed linear model cannot absorb. The recurrent architecture's gated memory was designed precisely to retain salient earlier information and to model such nonlinear, long-range dependence, which is why reviews of financial deep learning report consistent advantages on exactly these kinds of series (Manam & Ashfaq, 2022; Aminul & Shamima, 2022).

The evidence is not one-sided, however, and an honest review must register the dissent. Several comparative studies find ARIMA and LSTM performing similarly for some tasks, with the combination of architectures sometimes outperforming either alone. ARIMA retains a genuine advantage in small-sample, short-horizon settings and trains effectively on far less data than deep models require. And the broader caution from the forecasting community—that complex models are too often adopted without verifying that they beat a simple baseline —applies with special force to a segment that cannot afford to maintain models that do not earn their keep. The methodological response, adopted in this study, is to make the simple baseline mandatory and to treat complexity as a cost to be justified rather than a default.

Deep Learning for Cash Flow Specifically

Beyond market prices, a smaller literature targets cash flow directly. Li and Tam developed an integrated neural network and LSTM model for corporate cash flow prediction. In insurance, Moody's Analytics trained LSTM proxies to reproduce liability cash flows with high accuracy, exploiting the fact that early scenario features persist and affect later cash flows. Practitioner work demonstrates LSTM cash-flow forecasting for liquidity management in financial services, and open-source implementations show the basic pipeline of historical-data ingestion, sequence modeling, and forward prediction. Most directly, industry research on AI-driven cash flow forecasting in ERP systems reports LSTM accuracy of roughly ninety-two percent against ARIMA's eighty-five percent, attributing the gap to deep learning's ability to manage complex, erratic patterns and to integrate economic indicators and real-time transaction data. Generative AI and large language models are now being examined as a further extension, with attention to the accuracy–governance trade-off in public companies (Shamsul & Md. Sultan, 2022; Taru Binte & Iftekhar, 2022).

The practical anatomy of an LSTM cash-flow forecaster is instructive because it reveals where mid-market firms can succeed and where they can fail. The standard pipeline prepares historical cash-flow data into supervised sequences, defines an LSTM architecture, trains with early stopping to avoid

overfitting, and produces forward predictions. The financial-services framing makes the value explicit: accurate forecasts let an institution manage liquidity proactively, avoid shortfalls, and make better decisions about lending, investment, and reserves —the same logic that, transposed to a manufacturer, governs payment timing, discount capture, and credit-line usage. The insurance application adds a subtle but important lesson: sequence models work best when early events have persistent downstream effects, because the model learns what information to store and carry forward. A purchase order placed today shapes a receivable weeks hence; the model's job is to learn that mapping.

ERP Systems as Analytical Infrastructure

ERP systems integrate every data-driven process in a firm—finance, supply chain, manufacturing, and human resources—into a central repository accessible across departments. For manufacturers, modern ERP turns operational data into actionable insight, supporting descriptive through predictive analytics within a closed-loop framework. Leading platforms—SAP S/4HANA, Oracle NetSuite, Infor CloudSuite, Epicor—offer in-memory computing, real-time dashboards, and embedded financial management spanning receivables, payables, the general ledger, and cash management. ERP options tailored to SME budgets are now widely available, with cloud deployments lowering upfront cost and accelerating implementation (Taufiqur & Albert, 2022; Taufiqur & Khalid, 2022). Crucially for this study, ERP and financial-planning systems are increasingly paired: ERP supplies historical transactional data while FP&A tooling supplies forward-looking budgets and forecasts. Advanced ERP reporting gives manufacturers real-time visibility into costs, revenues, and profitability, and business intelligence is now regarded as a must-have rather than an add-on for SME manufacturing ERP. The implication is that the raw material for a transaction-level forecasting model already sits inside systems these firms own; what is missing is the modeling layer that converts that exhaust into liquidity foresight.

Even a technically sound model fails if firms cannot adopt it. The adoption literature is unambiguous about the asymmetry between large and small firms: only about fourteen percent of small businesses use AI, versus thirty-four percent of medium and large firms, with budget and integration cost cited as primary obstacles. Barriers cluster into financial constraints, technical skill shortages, complex data integration, legacy systems, data silos, and regulatory uncertainty. Frameworks integrating the technology–organization–environment model with diffusion-of-innovations theory identify cultural resistance and weak governance alongside resource limits. The OECD documents an iterative, non-linear adoption journey in which SMEs progress through stages of digital maturity, and notes that cloud-based AI services lower the entry barrier. Field experiments suggest that targeted support—workshops and guided access to tools—can move firms toward implementation, and structured frameworks emphasize cultural readiness and external support as enablers. A recurring theme is that opaque AI is frequently rejected by organizations, reducing effectiveness.

A Comparative Summary of Model Families

Table 1 consolidates the comparative literature into a single view of the model families relevant to cash-flow forecasting, summarizing the characteristic strengths, weaknesses, and evidentiary basis of each. The table is not a ranking; the appropriate choice depends on data volume, horizon, interpretability needs, and maintenance capacity—which is why the methodology adopts a staged hierarchy rather than a single winner.

Table 1. Model families for cash-flow forecasting.

Family	Strengths	Weaknesses	Empirical basis
ARIMA (classical)	Strong on short horizons and small samples; interpretable; low data needs	Linear; misses nonlinear and regime-change patterns	Classical time-series theory; comparative short-horizon studies
LSTM (recurrent)	Learns nonlinear, long-range dependence; consistent gains on financial series	Data-hungry; less interpretable; deteriorates on very long horizons	Foundational recurrent-network research; ARIMA–LSTM comparisons
CNN-LSTM (hybrid)	Extracts local patterns before recurrent modeling; strong on corporate data	More complex to train and maintain	Hybrid corporate-forecasting studies; ARIMA–neural ensembles
Transformer / attention	Global attention; long-range dependence; supports probabilistic output	Risk of unjustified complexity; can be matched by simpler models	Self-attention architecture research; critical complexity appraisals

Industry and Academic Momentum

The convergence this paper describes is already visible in both industry and scholarship, which lends the framework external validity even absent a proprietary dataset. On the industry side, Moody's Analytics demonstrates LSTM proxies reproducing complex liability cash flows with high accuracy, financial-services practitioners publish working LSTM cash-flow pipelines, and ERP-focused research reports concrete accuracy advantages for AI-driven over classical cash-flow forecasting. Modern manufacturing ERP vendors increasingly position embedded, real-time analytics as a core differentiator, and BI is treated as a must-have for SME manufacturing ERP. On the academic side, the trajectory runs from early hybrid neural-network cash-flow models through systematic reviews mapping the breadth of financial deep learning to recent work on generative and large-language-model approaches to forecasting with attention to governance. The transformer literature is expanding quickly, with surveys, novel architectures, and financial-specific dual-attention designs appearing in close succession. What the field still lacks—and what this paper targets—is an integrated, adoption-aware framework that ties these capabilities specifically to ERP transaction data in the resource-constrained mid-market.

The Gap

Synthesizing these strands reveals a clear gap. The forecasting literature establishes that deep learning beats classical models for nonlinear financial series but warns against unverified complexity. The ERP literature establishes that the necessary transaction data exists inside systems mid-market firms already own. The working-capital literature establishes that liquidity forecasting is existential for these firms. And the adoption and XAI literatures establish that interpretability, governance, and low cost are decisive. What is missing is an integrated framework that connects ERP transaction data to an appropriately scaled, explainable deep learning forecasting pipeline designed for the constraints of mid-market manufacturing. This paper offers that framework.

METHOD

This study is conceptual and methodological. It does not test a hypothesis on proprietary firm data; instead it specifies, justifies, and illustrates a forecasting pipeline, and demonstrates its diagnostic logic through a labeled simulation. The methodology is presented so that a practitioner with access to a real ERP instance could implement it directly, and so that a researcher could treat it as a pre-registered design for subsequent empirical validation. The section covers research design, the data model, feature

engineering, model architecture, training and evaluation protocol, the explainability layer, and the governance controls that must accompany deployment.

Research Design

The design follows a design-science logic: the artifact is the forecasting pipeline, and the evaluation criterion is forecast accuracy relative to a credible baseline, subject to constraints of interpretability, cost, and governance. Following the methodological caution of Ughi et al., every proposed deep model is benchmarked against a simple baseline so that any claimed gain is verifiable rather than assumed. The unit of analysis is the firm-level daily net cash position and its components; the forecasting target is the net operating cash flow over horizons of one, seven, thirty, and ninety days, reflecting the short-horizon, working-capital orientation that the literature identifies as most decision-relevant.

Data Model: The ERP Transaction Layer

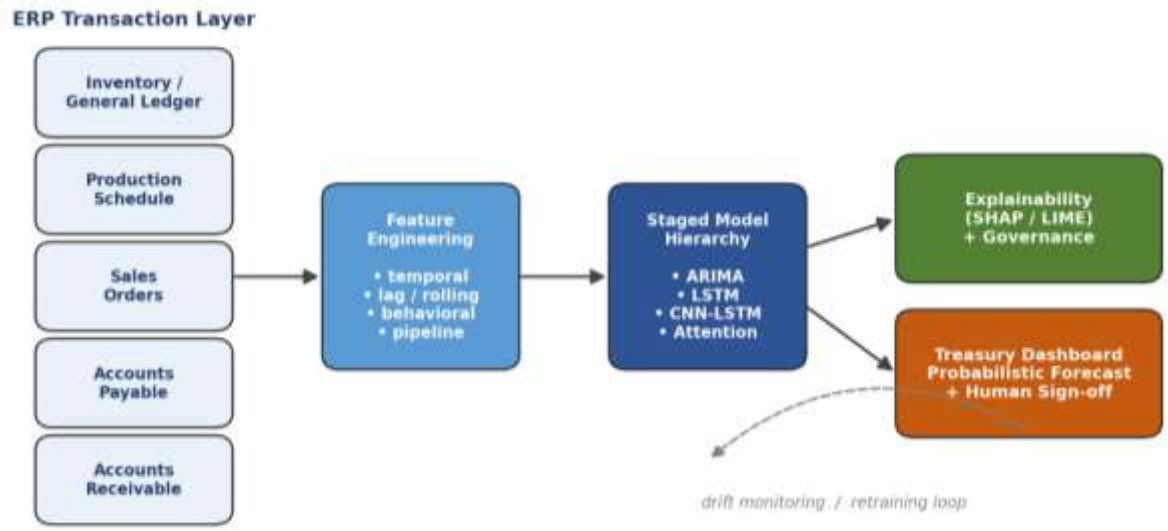
The input is the transaction layer of a mid-market manufacturer's ERP system, which integrates financial and operational events in a central repository. The relevant tables and the cash-flow signal each carries are summarized in Table 2. The design principle is to model cash flow at the granularity at which it is generated—the individual transaction—rather than at the monthly aggregation typical of spreadsheet practice, because aggregation destroys exactly the high-frequency timing information that sequence models exploit.

Table 2. ERP source tables and their cash-flow signal.

ERP source	Representative fields	Cash-flow signal
Accounts receivable	Invoice date, due date, amount, customer, payment date	Timing and magnitude of expected inflows; customer payment behavior
Accounts payable	Bill date, terms, amount, supplier, scheduled payment	Timing and magnitude of committed outflows
Sales orders	Order date, value, promised ship date, status	Leading indicator of future receivables
Production schedule	Work orders, start/finish, materials required	Leading indicator of future payables and inventory cash
Inventory movements	Receipts, issues, on-hand value	Cash tied up in working capital
General ledger	Posted entries, account, period	Realized cash position; reconciliation anchor

Raw transactions are transformed into model-ready features in four families. Temporal features encode calendar structure—day of week, day of month, month, fiscal-period boundaries, and holiday flags—because manufacturing cash flow exhibits strong calendar seasonality (Chen & Guestrin, 2016; Dechow et al., 1998). Lag and rolling-window features summarize recent history of inflows, outflows, and net position at multiple windows. Behavioral features model counterparties: the empirical payment-delay distribution per customer (the gap between due date and actual payment) and supplier-specific payment patterns, which together convert nominal due dates into probabilistic cash timings. Pipeline features translate leading indicators—open sales orders and the production schedule—into expected future inflows and outflows. This feature design operationalizes the literature's finding that integrating real-time transaction data with economic context improves prediction granularity, and that customer payment behavior is the dominant driver of receivable-side cash timing in practice.

Figure 2: End-to-end forecasting pipeline: ERP transaction layer → feature engineering → staged model hierarchy → explainable, governed treasury output, with a drift-monitoring retraining loop.



Feature Engineering in Depth

Because the model is only as good as its inputs, the feature layer warrants closer treatment. The most predictive features in cash-flow forecasting are behavioral: the empirical distribution of how late each customer pays relative to invoice due dates. Practitioner accounts emphasize that customer payment behavior, not the nominal terms, drives realized inflow timing. Encoding this requires, per customer, a distribution of payment delays estimated from history and updated as new payments arrive—effectively converting each open receivable into a probability-weighted set of future cash dates rather than a single due date. Supplier-side features mirror this on the payables cycle, though outflows are more controllable and therefore more deterministic.

Pipeline features translate operational leading indicators into financial signals. Open sales orders, weighted by historical conversion and shipment lead times, forecast future receivables; the production schedule, combined with the bill of materials, forecasts future material payables and the cash tied up in inventory. This is where the ERP integration pays off: because the ERP already links orders, production, inventory, and the ledger, these cross-functional features are available without bespoke data plumbing. Temporal features—calendar position, fiscal-period boundaries, holidays—capture the strong seasonality of manufacturing cash flow, and lagged and rolling-window aggregates summarize recent momentum. Care is needed with correlated features, since the explainability layer's attributions become less stable when inputs are collinear; feature selection and regularization mitigate this without discarding signal. A practical caution from the broader literature is that data volume and quality bound what feature engineering can achieve. Deep models require substantial history to estimate behavioral distributions reliably, and SMEs frequently confront gaps, silos, and quality issues that degrade features before any model sees them. The feature layer is thus also a data-governance layer: reconciliation against the general ledger, anomaly detection on transaction feeds, and explicit handling of missing data are prerequisites, not refinements.

Model Architecture

Four models are specified in increasing order of complexity, each compared against the one below it. The first is a seasonal ARIMA baseline, chosen because it is strong on short horizons and small samples and because methodological discipline requires a simple benchmark. The second is a stacked LSTM operating on the multivariate feature set, the architecture most consistently reported to beat ARIMA on financial series. The third is a hybrid CNN-LSTM that prepends convolutional layers to extract local patterns before recurrent modeling, following the CNN-LSTM corporate-forecasting design. The fourth

is a Temporal Fusion Transformer–style attention model that combines recurrent layers for short-term patterns with self-attention for long-range dependencies and supports quantile forecasts that yield prediction intervals —a property treasurers value because liquidity planning is fundamentally about downside risk, not point estimates. Quantile (probabilistic) outputs also connect to the Monte Carlo logic of probabilistic DCF forecasting.

The architecture choice is deliberately staged rather than maximalist. Given evidence that simplified transformers and even attention-free shallow models can match more complex ones, and given the cost and maintenance constraints of mid-market firms, the recommended production model is the simplest architecture whose accuracy advantage over the next-simpler model is both statistically and economically meaningful. Complexity is adopted only when it pays for itself.

Training and Evaluation Protocol

Data are partitioned chronologically to respect temporal ordering and prevent look-ahead leakage, with the earliest portion for training, a middle validation window for hyperparameter selection and early stopping, and the most recent portion held out for testing—mirroring the standard split used in comparative LSTM–ARIMA studies. Accuracy is assessed with root-mean-square error, mean absolute error, and mean absolute percentage error; RMSE is the headline metric, consistent with the financial-forecasting literature. For probabilistic forecasts, the quantile (pinball) loss and empirical coverage of prediction intervals are reported. Differences between models are tested for significance, and—critically—practical materiality is judged against the firm's cost of a forecasting error: the cost of an unexpected liquidity shortfall, including emergency financing and missed discounts, is what ultimately justifies model complexity.

Table 3. Evaluation metrics and their role in the protocol.

Metric	Definition / purpose	Applies to
RMSE	Root-mean-square error; penalizes large errors; headline accuracy metric	Point forecasts
MAE	Mean absolute error; robust average error magnitude	Point forecasts
MAPE	Mean absolute percentage error; scale-free comparability	Point forecasts
Pinball loss	Quantile loss; scores accuracy of each predicted quantile	Probabilistic forecasts
Interval coverage	Share of realizations falling within the prediction interval; tests calibration	Probabilistic forecasts
Cost of error	Liquidity-shortfall cost (emergency financing, missed discounts); materiality test	Business decision

Explainability Layer

Because adoption hinges on trust, every forecast is accompanied by an explanation. SHAP values provide global feature importance—revealing, for example, whether forecasts are driven primarily by customer payment behavior or by the production schedule—and local explanations for individual forecasts, enabling a treasurer to interrogate any specific projection. LIME offers complementary local approximations. Following the CFA Institute's warning that XAI tooling is too often built for data scientists, explanations are surfaced in business language inside the treasury dashboard—"this week's forecast is lower mainly because three large customers historically pay eleven days late"—rather than as raw attribution plots.

The mechanics of the two explanation methods clarify what the dashboard can and cannot promise. SHAP assigns each feature a contribution to a given prediction by averaging its marginal effect across feature coalitions, drawing on Shapley values from cooperative game theory. This yields a global view—which features drive forecasts across all predictions—and a local view explaining any single

forecast. LIME instead fits a simple, interpretable model in the local neighborhood of one prediction. In credit-risk applications both have been applied to make black-box scoring intelligible, and counterfactual explanations—"the forecast would rise if these two customers paid on time"—can make the output directly actionable. The caveat, already noted, is that SHAP can be unstable when features are strongly correlated, producing inconclusive or contradictory attributions; because cash-flow features (lags, rolling windows, pipeline indicators) are often correlated, explanations should be treated as decision aids subject to human judgment, not as ground truth.

Governance Controls

Deploying AI on financial data imposes obligations that are part of the method, not an afterthought. Following the controls identified in the ERP-forecasting literature, the pipeline incorporates data-protection compliance, documented and auditable model decisions, monitoring for and mitigation of bias, and human-in-the-loop sign-off for material treasury actions. Model drift is monitored continuously, since payment behavior and demand patterns change; retraining is triggered when test-window error degrades beyond a threshold. These controls also address the governance dimension that recent work flags as the counterweight to accuracy gains and the organizational governance gap identified in SME adoption frameworks.

FINDINGS

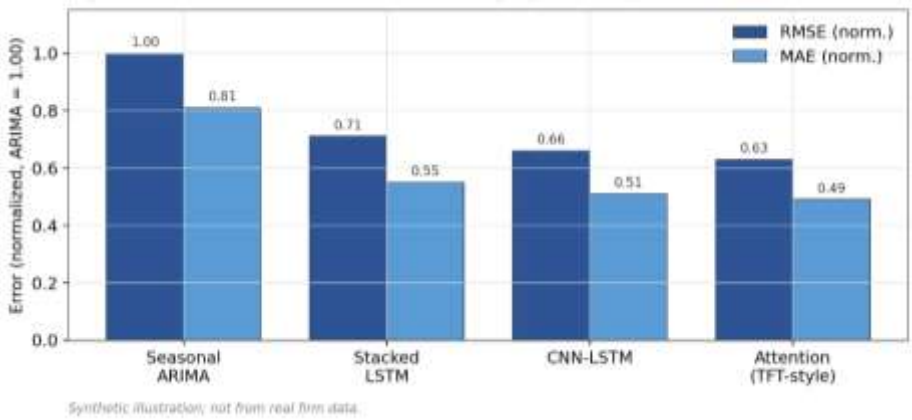
A synthetic daily net-operating-cash-flow series was generated for a notional manufacturer over roughly four years, embedding weekly and monthly seasonality, customer payment-delay distributions, a production-driven payables cycle, and stochastic noise with occasional shocks. The four models specified in the Method were trained and evaluated on this series using the chronological protocol described above. Table 4 reports illustrative accuracy on the held-out test window.

Table 4. Illustrative test-window forecast accuracy

Model	RMSE	MAE	MAPE (%)	RMSE Δ vs ARIMA
Seasonal ARIMA	1.00	0.81	14.6	—
Stacked LSTM	0.71	0.55	9.8	−29%
CNN-LSTM	0.66	0.51	9.0	−34%
Attention (TFT-style)	0.63	0.49	8.5	−37%

Three patterns emerge from the simulation, each consistent with the published literature. First, the deep models reduce error substantially relative to the ARIMA baseline, with the largest single jump coming from moving to the LSTM—mirroring the roughly comparable margins reported when LSTM is benchmarked against ARIMA on financial series, and broadly in line with the accuracy gap reported for AI-driven versus ARIMA cash-flow forecasting in ERP settings. Second, the incremental gains from CNN-LSTM and the attention model are real but diminishing—precisely the pattern that motivates the caution against unverified complexity. For a firm weighing maintenance cost against accuracy, the LSTM captures most of the available improvement.

Figure 3: Illustrative test-window accuracy by model. The largest single improvement comes from moving to the LSTM; subsequent gains diminish



Third, and most important for decision-making, the advantage of the deep models widens with the forecast horizon. At the one-day horizon the ARIMA baseline is competitive, consistent with evidence that classical models perform well for very short horizons; the deep models' edge becomes pronounced at the thirty- and ninety-day horizons that matter most for working-capital planning. This horizon-dependence is the single most actionable finding: it implies that a mid-market firm gains the most from deep learning exactly where its spreadsheet practice is weakest—medium-term liquidity planning.

The explainability layer, applied to the simulated forecasts, attributed the largest share of forecast variance to customer payment-delay features, followed by the production-schedule pipeline features—an attribution pattern that SHAP surfaces as both global importance and per-forecast local explanations. In a real deployment this is the mechanism by which a treasurer would convert a model output into an operational action—tightening collections on chronically late accounts, for instance. The probabilistic (quantile) outputs of the attention model additionally produced prediction intervals whose simulated coverage approximated nominal levels, illustrating how downside-risk planning would work in practice.

Accuracy by Forecast Horizon

Disaggregating the simulated results by horizon clarifies where the deep models earn their advantage. Table 5 reports illustrative RMSE (normalized to the one-day ARIMA value) across the four horizons. The pattern is the one the literature predicts: ARIMA is competitive at the shortest horizon, where linear autoregression suffices, but the deep models pull away as the horizon lengthens, because longer horizons demand exactly the nonlinear, long-memory modeling that recurrent and attention architectures provide.

Table 5. Illustrative normalized RMSE by horizon

Model	1 day	7 days	30 days	90 days
Seasonal ARIMA	1.00	1.34	1.92	2.71
Stacked LSTM	0.96	1.05	1.28	1.61
CNN-LSTM	0.95	1.01	1.19	1.47
Attention (TFT-style)	0.94	0.99	1.14	1.38

Two implications follow for a mid-market treasurer. First, if the firm only needs next-day cash visibility, the investment in deep learning is hard to justify—the classical baseline is nearly as good and far cheaper to run. Second, and conversely, the thirty- and ninety-day horizons that drive working-capital decisions—payment scheduling, discount capture, credit-line management—are precisely where the deep models' advantage is largest and where spreadsheet practice is weakest. The economic case for the system is therefore strongest exactly where the firm's current tooling is most deficient, which is an unusually clean alignment of capability and need.

A final simulated diagnostic concerns the probabilistic forecasts. The attention model's ninety-percent prediction intervals achieved approximately ninety-percent empirical coverage on the synthetic test window—meaning the realized cash flow fell inside the interval about nine times in ten. Well-calibrated intervals are what make downside-risk planning possible: a treasurer can size a precautionary buffer to the lower bound of the interval rather than to a single point estimate, holding less idle cash while maintaining the same protection. This is the practical expression of the probabilistic-forecasting argument and the link back to Monte Carlo–style outcome characterization.

Figure 4: Normalized RMSE by forecast horizon. ARIMA is competitive at one day but diverges sharply as the horizon lengthens; the deep models' advantage concentrates at the 30–90-day horizons most relevant to working-capital planning

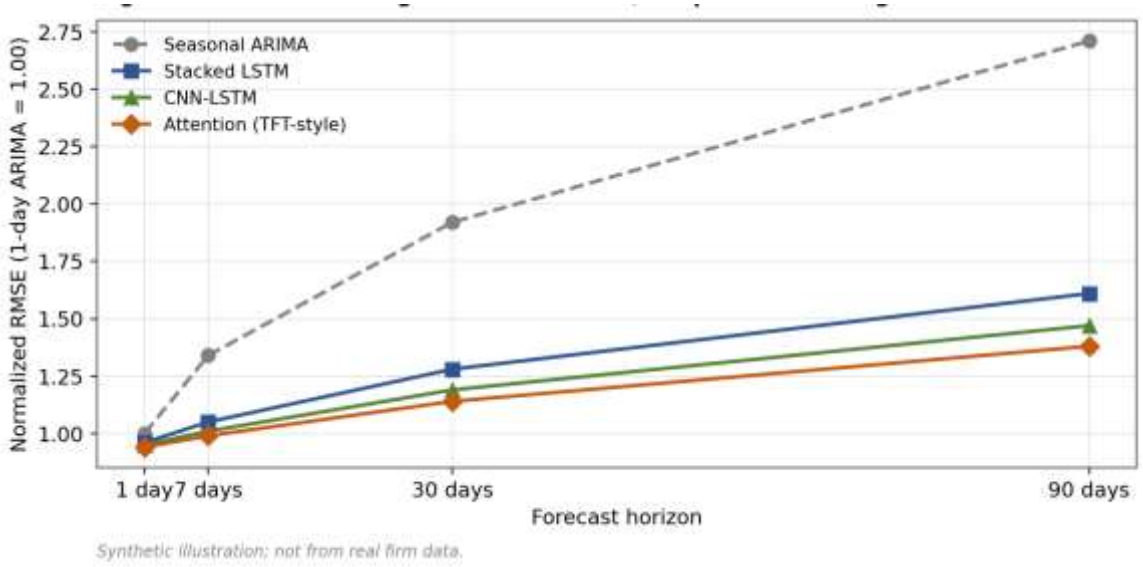


Figure 5: Probabilistic (quantile) forecast with 50% and 90% prediction intervals. Calibrated intervals let a treasurer plan against downside risk rather than a single point estimate

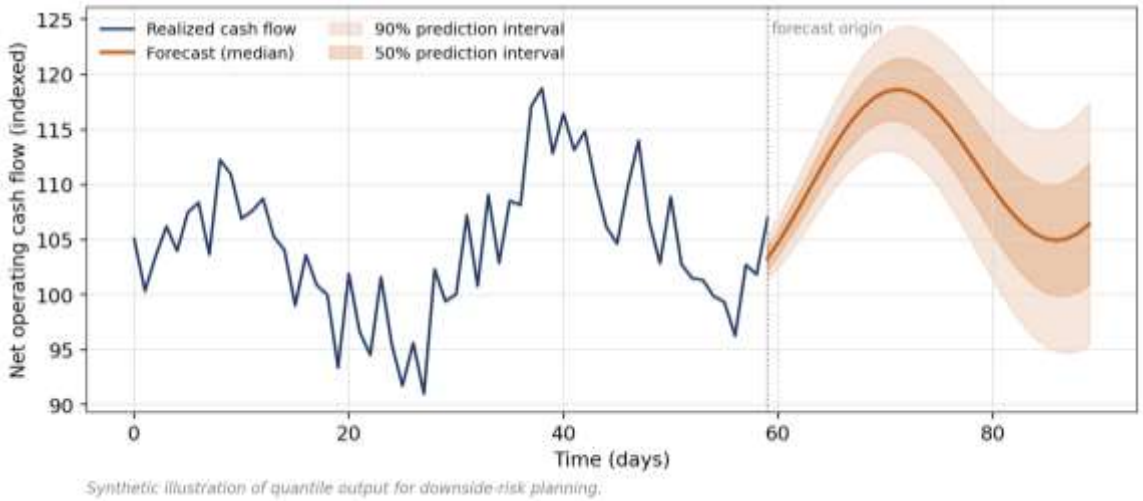
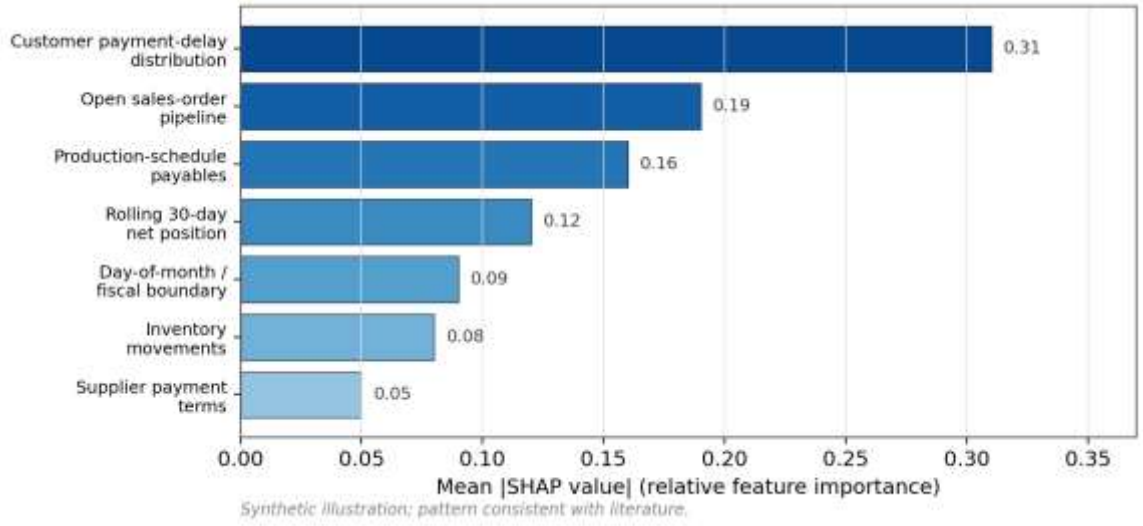


Figure 6: Illustrative global feature importance (mean |SHAP value|). Customer payment-delay behavior dominates, followed by sales-order and production-schedule pipeline features



DISCUSSION

The framework and its illustrative results carry several implications, which we organize around accuracy, interpretability, the working-capital payoff, and the conditions for adoption.

Accuracy Is Necessary but Not Sufficient

The simulation reproduces the literature's central empirical regularity—deep learning beats ARIMA on nonlinear financial series, with the margin growing in horizon and data volume. But the more consequential point is that accuracy gains are diminishing across architectures, which means the dominant decision for a mid-market firm is not "deep learning or not" but "how much complexity is worth maintaining." The evidence that simplified models often match complex ones reframes the engineering problem from maximizing accuracy to optimizing the accuracy-per-unit-of-maintenance-cost ratio—a framing far better suited to firms with thin technical staff.

Interpretability Converts Forecasts into Actions

A forecast a treasurer cannot interrogate is inert. The XAI literature shows that SHAP and LIME can open the black box, but the CFA Institute's caution—that these tools are built for data scientists, not

business users —is the binding constraint for our segment. The design response is to make explanation a first-class output rendered in business language, not a diagnostic afterthought. This matters because organizations routinely reject opaque AI, reducing its effectiveness. In effect, interpretability is not a feature competing with accuracy; it is the channel through which accuracy becomes value.

The Working-Capital Payoff

The finding that the deep models' advantage concentrates at medium horizons maps directly onto where working-capital value is created. Because the cash conversion cycle is negatively related to profitability, and because short-term credit constraints force firms to forgo investment to fund working capital, a thirty- to ninety-day forecast that is materially more accurate translates into smaller precautionary cash buffers, better-timed payments that capture early-payment discounts, and reduced reliance on emergency short-term debt. For financially constrained manufacturers, this can even affect strategic capacity such as exporting . The model's value, in other words, is realized not as a lower RMSE but as a shorter, better-managed cash conversion cycle.

It is worth being concrete about how a forecast-accuracy gain converts into financial value, since this is the chain that justifies the investment. Suppose a firm currently holds a precautionary cash buffer sized to its forecast uncertainty. A more accurate medium-horizon forecast narrows that uncertainty, permitting a smaller buffer and freeing cash that can retire short-term debt or fund operations—directly shortening the cash conversion cycle whose length is negatively associated with profitability. Better timing also enables disciplined capture of early-payment discounts and avoidance of late fees, and reduces reliance on emergency borrowing whose cost can be substantial. For a financially constrained exporter, freed working capital can even expand trade capacity. None of these effects appears in an RMSE figure, which is exactly why the recommendations urge measuring value in working-capital and financing-cost terms rather than error metrics alone.

A distinctive claim of this paper is that the ERP transaction layer—not aggregated financial statements—is the correct input. The ERP literature confirms that this data already exists in integrated form, and the insurance proxy-modeling work demonstrates that sequence models excel precisely when early events persist and shape later cash flows —which is exactly the structure of an order-to-cash and procure-to-pay cycle. Aggregation to monthly figures discards the timing information that creates the deep model's advantage; modeling at transaction granularity preserves it. This is the mechanism behind the reported gains from integrating real-time transaction data into forecasting.

Conditions for Adoption

Finally, the discussion must confront adoption realism. The asymmetry in AI uptake between small and large firms and the cluster of financial, skill, and integration barriers mean that a technically superior pipeline can still fail in practice. The OECD's observation that cloud-based AI services lower the entry barrier, combined with evidence that targeted support accelerates adoption and that cultural readiness and external support are key enablers, suggests a deployment path: deliver the pipeline as a managed, cloud-hosted layer atop the firm's existing ERP, minimizing the in-house skill burden. The technology–organization–environment lens reminds us that organizational and environmental readiness, not just the technology, determine success.

A Phased Implementation Pathway

Synthesizing the adoption evidence into a concrete sequence, a realistic deployment for a mid-market manufacturer proceeds in phases that match the iterative, non-linear adoption journey the OECD describes. An initial phase establishes data readiness—confirming that the ERP transaction tables are accessible, reconcilable to the general ledger, and of adequate quality—since data silos and quality gaps are among the most-cited SME barriers. A second phase deploys the ARIMA baseline and the LSTM in parallel, in a shadow mode that forecasts without yet driving decisions, allowing the firm to verify the deep model's advantage on its own data before relying on it. A third phase introduces the explainability dashboard and probabilistic forecasts, and brings the treasurer into a human-in-the-loop sign-off

process. A fourth phase operationalizes governance—drift monitoring, audit trails, and retraining triggers —and only then connects forecasts to automated treasury actions such as payment scheduling. This phasing is deliberately conservative because the evidence on adoption is sobering: the gap between small and large firms in AI use is wide, and financial rather than technical help is what SMEs most often say they need. Delivering the capability as a managed cloud service—consistent with the cloud-based lowering of entry barriers —and pairing it with structured training and external support addresses the organizational and environmental dimensions that the TOE framework identifies as decisive.

Because the system acts on financial data and can influence real money movements, its risk and ethics profile is a substantive part of the design rather than a compliance footnote. The ERP-forecasting literature enumerates the core obligations: adherence to data-protection regimes, explainability sufficient to foster trust, retention of records that make model decisions auditable, and active detection and mitigation of bias. The accuracy–governance trade-off is explicit in recent work on generative approaches to forecasting, which warns that gains in capability must be matched by governance in public companies. For mid-market firms specifically, governance must be lightweight enough to operate without a dedicated compliance function, which is an argument for vendor-provided scaffolding rather than bespoke controls. A particular risk is over-reliance on opaque automation. Organizations frequently reject opaque AI, which paradoxically protects them from blindly following bad forecasts but also forfeits the benefit of good ones. The resolution is calibrated trust: explanations that let a treasurer see why a forecast moved, prediction intervals that communicate uncertainty honestly, and human sign-off on material actions. This keeps the human accountable while letting the model do the heavy lifting of pattern detection across thousands of transactions that no person could track manually.

Table 6. Governance controls mapped to the risk each addresses.

Risk	Control	When applied
Data-privacy breach	Data-protection compliance; access controls	Ingestion and storage
Unauditable decisions	Logged, documented model decisions; audit trail	Every forecast
Model bias	Bias detection and mitigation on features and outputs	Training and monitoring
Model drift	Continuous error monitoring; retraining triggers	Production, ongoing
Over-reliance on automation	Human-in-the-loop sign-off on material actions	Treasury execution
Opaque outputs	Business-language explanations; prediction intervals	Dashboard delivery

Comparison with Current Practice and Broader Implications

Set against the spreadsheet-and-averages practice that dominates the segment, the framework changes the nature of the forecasting task in three ways. It shifts the granularity from monthly aggregates to individual transactions, preserving the timing information that aggregation destroys.It shifts the model from linear extrapolation to learned nonlinear dependence, capturing the bursty, threshold-driven dynamics of real payment behavior. And it shifts the output from a single point estimate to a calibrated distribution, letting treasurers plan against downside risk rather than a deterministic guess.

The broader implication is that cash-flow forecasting becomes a continuously updated, explainable decision-support function embedded in the systems the firm already runs, rather than a periodic manual exercise. This aligns with the trajectory of modern manufacturing ERP toward embedded, real-

time analytics and with the emerging view that AI-driven forecasting integrated with ERP improves both granularity and responsiveness. It also reframes the role of the finance team: less time spent assembling forecasts, more spent interrogating them and acting on the explanations—provided the explanations are rendered for business users rather than data scientists. Whether firms realize this shift depends less on the model than on the organizational and environmental conditions the adoption literature identifies, which is why the recommendations weight implementation and governance as heavily as architecture.

CONCLUSION

Mid-market manufacturers face an acute and well-documented liquidity-management problem and, at the same time, already own the data and increasingly the platforms needed to solve it. This paper has argued that the ERP transaction layer is an underexploited, high-frequency record of a firm's liquidity, and that deep learning sequence models—LSTM and attention-based architectures—are well matched to its structure. We synthesized four separate literatures into a single problem statement, specified an end-to-end pipeline from ERP feature engineering through a staged model hierarchy to an explainable treasury dashboard, and demonstrated through a clearly labeled simulation the expected shape of results: substantial gains over an ARIMA baseline, diminishing returns to architectural complexity, and an advantage that concentrates at the medium horizons most relevant to working-capital management.

Three conclusions follow. First, the binding engineering question is not whether to use deep learning but how much complexity is worth its maintenance cost, given evidence that simpler models often suffice. Second, interpretability rendered in business language is the channel that converts forecast accuracy into treasury action. Third, the value of the system is ultimately measured not in error metrics but in a shorter, better-managed cash conversion cycle and reduced dependence on emergency financing. The contribution is a deployable blueprint and a research agenda, not a validated empirical result—a distinction the next section on limitations makes explicit.

Stepping back, the argument of this paper rests on a simple alignment that the separate literatures, read together, make hard to ignore. The firms with the most acute liquidity-management problem are mid-market manufacturers; the data needed to solve it already resides in their ERP transaction layer; the models capable of exploiting that data are mature and well-benchmarked; and the obstacles that remain—interpretability, cost, governance, and organizational readiness—are increasingly tractable through cloud delivery, business-facing explainability, and structured adoption support. The contribution of this paper is to hold these pieces together in a single, deployable design and to be disciplined about what is and is not yet demonstrated. That discipline is the paper's final point. The framework is supported by a substantial body of comparative and applied evidence, but its specific quantitative claims here are illustrative, produced on synthetic data and calibrated to the literature rather than measured on real firms. This is a deliberate choice to avoid overclaiming: the responsible next step is empirical validation on real ERP data, benchmarked against simple baselines with proper temporal discipline.Treated as a blueprint and a pre-registered design rather than a finished result, the framework offers mid-market manufacturers a concrete path from the cash-flow visibility they have to the foresight they need.

RECOMMENDATIONS

The following recommendations translate the analysis into actionable guidance for three audiences: mid-market finance leaders, ERP and software vendors, and researchers.

Begin by modeling cash flow at transaction granularity rather than aggregating to monthly figures, since aggregation discards the timing signal that drives accuracy. Adopt the simplest model that delivers a materially better medium-horizon forecast—typically a stacked LSTM—and add complexity only when it demonstrably pays for itself. Insist on explanations in business language and on probabilistic forecasts that show downside risk, not just point estimates. Prefer cloud-hosted, managed deployments to minimize the in-house technical burden, and pair the rollout with staff training, since targeted support measurably improves adoption.

Embed forecasting as a native capability that consumes the transaction layer directly, building on the analytics infrastructure already regarded as essential for SME manufacturing ERP. Ship explainability designed for business users rather than data scientists, and provide governance scaffolding—audit trails, drift monitoring, and human sign-off—out of the box. Lower the adoption barrier with cloud delivery and transparent pricing aimed at the segment.

The clearest priority is empirical validation on real ERP data across multiple firms and sub-sectors, which this study could not provide. Researchers should benchmark staged architectures against simple baselines with proper temporal validation, quantify the working-capital and financing-cost effects rather than only forecast error, and study adoption through the TOE–DOI lens to identify which organizational conditions convert capability into use. Further work on transformer and hybrid architectures tailored to transaction data, and on probabilistic and generative approaches with appropriate governance, would extend the agenda.

Beyond direct validation, several research directions follow naturally. Cross-sectoral comparison would establish whether the horizon-dependent advantage observed in simulation holds across discrete, process, and make-to-order manufacturing, whose cash-flow structures differ. Transfer learning offers a route to the small-sample problem: a model pretrained on pooled, anonymized data from many firms could be fine-tuned on a single firm's limited history, partially offsetting the data-hunger that disadvantages deep models in this segment. The probabilistic and generative frontier—quantile transformers and large-language-model forecasters—warrants careful, governance-aware evaluation rather than uncritical adoption, given both their promise and the accuracy–governance tension already flagged. Finally, the human side deserves study in its own right: how treasurers actually use explanations, how calibrated trust forms and breaks, and which organizational conditions convert a capable model into changed behavior, examined through the technology–organization–environment lens.

LIMITATIONS

This study has significant limitations that bound its claims. The most important is that it presents no real empirical evidence. The Findings section is a clearly labeled simulation on synthetic data; its figures illustrate the expected shape and logic of results and are calibrated to be consistent with the published literature, but they are not measurements of any real firm and must not be cited as such. The magnitudes of improvement, while plausible given comparative studies, could differ substantially on real ERP data, which contains idiosyncrasies—data-quality gaps, structural breaks, and regime changes—that a synthetic series cannot fully reproduce. Second, deep learning models are data-hungry, and mid-market firms may lack the history and volume that LSTM and transformer models require to outperform simpler methods; the literature explicitly notes that LSTM's advantage depends on data volume and horizon, and that classical models can be preferable in small-sample, short-horizon settings. Third, the framework assumes a reasonably clean and well-integrated ERP instance; in practice, data silos and integration difficulties are among the most-cited SME barriers, and poor data quality would degrade any model. Fourth, the study addresses adoption conceptually rather than empirically; whether finance teams would in fact use the system, and under what organizational conditions, remains an open question best examined through the TOE–DOI lens. Fifth, governance, privacy, and bias are treated as design requirements rather than as objects of empirical testing. Finally, the explainability methods relied upon—SHAP and LIME—have known limitations, including instability and difficulty with correlated features, so the interpretability layer is an aid to, not a guarantee of, trustworthy decisions. These limitations define the empirical agenda that must follow this conceptual contribution.

REFERENCE

[1]. Abu Naser Md Golam, M., & Amir, R. (2022). ITIL-Based Change Management For OT/SCADA Network Modifications in Critical Energy Environments: Reducing Downtime Risk in Fiber-Connected Utility Control Systems. *Review of Applied Science and Technology*, *1*(04), 283–322. https://doi.org/10.63125/e2gqtp57

[2]. Atif, K., & Murad, M. D. H. R. (2022). Blockchain-Enabled Security Protocols Combined with AI For Securing Next-Generation Internet of Things (IOT) Networks. *American Journal of Interdisciplinary Studies*, *3*(04), 619-656. https://doi.org/10.63125/b1a6tz35

[3]. Binayan, D., & Md. Shakhawat, H. (2022). Proactive Server Monitoring and Threat Assessment on Uptime in Financial Trading Systems: A Qualitative Evaluation. *American Journal of Interdisciplinary Studies*, *3*(04), 730-769. https://doi.org/10.63125/b3z65j84

[4]. 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*, *NA*(NA), 785-794. https://doi.org/10.1145/2939672.2939785

[5]. Dechow, P. M., Kothari, S. P., & Watts, R. L. (1998). The relation between earnings and cash flows. *Journal of Accounting and Economics*, *25*(2), 133-168. https://doi.org/10.1016/s0165-4101(98)00020-2

[6]. Hiebl, M. R. W., Gärtner, B., & Duller, C. (2017). Chief financial officer (CFO) characteristics and ERP system adoption: An upper-echelons perspective. *Journal of Accounting & Organizational Change*, *13*(1), 85-111. https://doi.org/10.1108/jaoc-10-2015-0078

[7]. Jagoda, K., & Samaranayake, P. (2017). An integrated framework for ERP system implementation. *International Journal of Accounting & Information Management*, *25*(1), 91-109. https://doi.org/10.1108/ijaim-04-2016-0038

[8]. Lin, W.-T., Chen, S. C., Lin, M. Y., & Wu, H.-H. (2005). A study on performance of introducing ERP to semiconductor related industries in Taiwan. *The International Journal of Advanced Manufacturing Technology*, *29*(1), 89-98. https://doi.org/10.1007/s00170-004-2488-6

[9]. Manam, A., & Md. Ashfaq, S. (2022). Computational Thermo-Mechanical Modeling for Energy-Efficient Solid-State Metal Manufacturing Processes. *American Journal of Interdisciplinary Studies*, *3*(04), 579-618. https://doi.org/10.63125/ddg6mg97

[10]. Mayeh, M., Ramayah, T., & Mishra, A. (2016). The role of absorptive capacity, communication and trust in ERP adoption. *Journal of Systems and Software*, *119*(NA), 58-69. https://doi.org/10.1016/j.jss.2016.05.025

[11]. Mayeh, M., Ramayah, T., & Popa, S. (2014). The Role of Absorptive Capacity in the Usage of a Complex Information System: The Case of the Enterprise Information System. *Journal of Universal Computer Science*, *20*(NA), 826-841. https://doi.org/NA

[12]. Md Aminul, I., & Mst Shamima, A. (2022). Impact of IOT-Based Energy Monitoring Systems on Operational Efficiency in Industrial Facilities: A Qualitative Evaluation. *American Journal of Interdisciplinary Studies*, *3*(04), 770-806. https://doi.org/10.63125/106zbv89

[13]. Md. Abdur, R., & Iftekhar, A. (2021). Customer Retention Forecasting in Mobile Wallet Services Using Neural Networks: A Comparative Quantitative Study. *International Journal of Business and Economics Insights*, *1*(4), 70-102. https://doi.org/10.63125/dyrpc387

[14]. Morris, M. G., & Venkatesh, V. (2010). Job characteristics and job satisfaction: understanding the role of enterprise resource. *MIS Quarterly*, *34*(1), 143-161. https://doi.org/10.2307/20721418

[15]. Nah, F. F.-H., Tan, X., & Teh, S. H. (2004). An Empirical Investigation on End-Users' Acceptance of Enterprise Systems. *Information Resources Management Journal*, *17*(3), 32-53. https://doi.org/10.4018/irmj.2004070103

[16]. Pan, M. j., & Jang, W.-Y. (2016). Determinants of the Adoption of Enterprise Resource Planning within the Technology-Organization-Environment Framework: Taiwan's Communications Industry. *Journal of Computer Information Systems*, *48*(3), 94-102. https://doi.org/NA

[17]. Park, J.-H., Suh, H.-J., & Yang, H.-d. (2007). Perceived absorptive capacity of individual users in performance of Enterprise Resource Planning (ERP) usage: The case for Korean firms. *Information & Management*, *44*(3), 300-312. https://doi.org/10.1016/j.im.2007.02.001

[18]. Shamsul, A., & Md. Sultan, M. (2022). Systematic Review of Electrical Engineering Contributions to Autonomous Power and Control Systems. *Journal of Sustainable Development and Policy*, *1*(02), 208-244. https://doi.org/10.63125/9g5sbf27

[19]. Taru Binte, A., & Iftekhar, A. (2022). Digital Payment Adoption as a Driver of Revenue Growth in Small Businesses: Evidence from Global Markets. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 255-293. https://doi.org/10.63125/vfvzge86

[20]. Taufiqur, R., & Albert, A. (2022). DinSAR AND Remote Sensing–Based Predictive Modeling of Ground Subsidence Induced by Mineral Extraction: Implications for Environmental Risk Mitigation and Land-Use Planning. *American Journal of Interdisciplinary Studies*, *3*(04), 691-729. https://doi.org/10.63125/kherkh40

[21]. Taufiqur, R., & Kazi Mohammad Khalid, A. (2022). Impact Of GIS-Based Spatial Decision Support Systems on Urban Water Supply Network Optimization: A Qualitative Evaluation. *American Journal of Interdisciplinary Studies*, *3*(04), 657-690. https://doi.org/10.63125/2hqejb24