

Playbook

From Fragmented Forecasts to AI-Guided Liquidity Decisions

A practical four-stage framework for modernizing forecasting, improving liquidity visibility, and preparing treasury for AI-enabled decision support.

Edouard Gabreau
SVP Product Management

Dory Malouf
Senior Director, Global Business Value Advisory

Michelle Woolfolk
Director Product Marketing



Contents

The Forecasting Efficiency Trap	The Four-Stage Forecasting Framework	Understanding Direct and Indirect Forecasting	Stage 1: Fragmented Operations	Stage 2: Connected Visibility
4	6	9	13	16
Stage 3: Intelligent Forecasting	Stage 4: AI Guided Liquidity Decisions	A Practical Maturity Path for Treasury Teams	The Liquidity Intelligence Gap Is Not a Technology Gap	
20	25	29	31	

“The biggest risk to your liquidity position isn’t market volatility or late payments.

It’s a forecasting process that can’t keep up with the speed of your business.”

This playbook is for treasury and finance leaders who know their forecast process needs to mature, but need a practical way to define what “better” looks like. It introduces a four-stage framework for moving from fragmented, spreadsheet-driven forecasting to connected visibility, intelligent forecasting, and AI-guided liquidity decisioning. Along the way, it shows what each stage enables, what capabilities matter most, and how to identify the next step without trying to leap straight to AI.



The Forecasting Efficiency Trap

Most treasury teams are working hard. They're pulling data from bank portals, reconciling ERP extracts, chasing inputs from business units, and carefully updating spreadsheets that took years to build. The problem isn't effort. The problem is that all of that work is going into assembling the forecast, leaving almost no time to actually use it. When the CFO needs an answer, the team is still building the picture.

A forecast can balance. It can be updated every week. It can even be directionally correct. But if it takes too long to produce, misses timing risk, hides category-level errors, or can't explain why cash moved, it's not giving the business what it needs.

The root cause is almost always the same: the information treasury needs is scattered across bank portals, ERP systems, regional spreadsheets, and email inboxes that don't talk to each other. By the time the forecast is ready, the moment to act has often already shifted. The challenge isn't talent or intent. It's infrastructure.



Customer PROOF

From scattered portals to one source of truth

"For Inchcape, the a-ha moment was seeing every bank account in one place. Instead of portals and spreadsheets scattered everywhere, we finally had a single source of truth, and I could trust Kyriba to look after all of our data."

Anthony Ho

Group Treasurer, Inchcape Shipping Services

[Read more](#)

Success in forecasting isn't a technology milestone. It's a behavioral and operational one.



Dory Malouf

"One of the biggest mistakes teams make is assuming forecast maturity starts with the most advanced tool. It actually starts with trust. If the CFO does not trust the inputs, the assumptions, or the process, the forecast will not change decisions, no matter how sophisticated the model is."

The real shift happens when treasury stops spending most of its time building the forecast and starts spending it using the forecast. A 10-day view becomes a 13-week rolling view. A spreadsheet that balances becomes an early warning system. A process that depends on specific people becomes one the whole team can trust.

The companies that will lead on liquidity over the next five years aren't necessarily the ones that invest in the most advanced technology first. They're the ones that build the right foundation and move through the stages in the right order.

A Point of View Worth Debating

Here is an opinion some people in this industry disagree with: the biggest predictor of forecasting maturity is not which tools your team has bought. It is whether your CFO is willing to act on an uncomfortable forecast. We have seen organizations with sophisticated technology that still smooth every number before it reaches the executive suite, compressing ranges and removing downside scenarios before the board meeting.

We have also seen teams on relatively simple platforms whose forecasts actually change decisions, because leadership trusts them enough to act. Technology enables that trust. It cannot create it. If your organization shoots the messenger, no amount of AI will fix your forecasting problem.

This playbook presents a practical four-stage framework to help treasury and finance leaders assess current capabilities, define optimal outcomes, and implement improvements through actionable, incremental steps.



The Four-Stage Forecasting Framework

Modern forecasting isn't a leap into the future. It's a step-by-step journey, where every capability built today becomes the floor for what you'll do tomorrow.

Think about how we navigate in the real world. We used to print MapQuest directions or struggle with paper maps. Then GPS showed up and told us where we actually were. Then it learned our patterns: how long routes usually took, which roads we preferred, where delays usually hit. Now your phone knows your destination, models three different routes, checks live traffic, and reroutes you before you even see the jam.

That's exactly how your forecasting should work. Stage 1 is the paper map era. Stage 2 gives you a trusted current position. Stage 3 adds prediction and planning. Stage 4 is a system that monitors, reasons, and has a recommendation ready before you've asked the question. The order matters: too many teams try to jump to Stage 4 insights while their data is still stuck in a mess of spreadsheets and inboxes. You can't get smart recommendations from broken inputs.



The Stage Framework

STAGE	MATURITY LEVEL	BUSINESS CAPABILITY	TECHNICAL CAPABILITY	CORE QUESTION	TREASURY SHIFT
Stage 1	Ad Hoc	Fragmented Operations	Manual and Rules-of-Thumb	What happened, and where did the money go?	From financial firefighter to structured data collector
Stage 2	Emerging	Connected Visibility	Connected Automation	Where is the cash right now, and can we trust that answer?	From data collector to data reviewer
Stage 3	Strategic	Intelligent Forecasting	Predictive Analytics and Strategic Modeling	Where will the cash be, what could change that, and how do we plan for it?	From data reviewer to decision advisor
Stage 4	Leader	AI-Guided Liquidity Decisioning	Agentic AI	What should treasury do next — and why?	From decision advisor to strategic governor with a human always in the loop



What each stage means for your forecast:

	STAGE 1	STAGE 2	STAGE 3	STAGE 4
Reliable Forecast Horizon	0–2 weeks, if tracked at all	Up to 13 weeks	3–18 months rolling	12–24 months, continuously updated
Time to Answer a Liquidity Question	Days; requires manual assembly	Hours; data is connected, and analysis is manual	Minutes; system surfaces actuals, scenarios, and variance, while humans interpret and decide	Proactive; the answer exists before the question is asked
Decisions Treasury Can Support	Explaining what already happened	Cash positioning; short-term funding calls	Predictive Analytics and Strategic Modeling	Capital structure optimization; M&A liquidity due diligence; investor guidance; balance sheet strategy
Forecast Confidence	Unknown — accuracy untracked	Moderate — current position trusted; future uncertain	Building to High — accuracy tracked by category; scenarios replace single-point estimates	Continuously calibrated; the system tracks its own accuracy and improves with every cycle

These four dimensions, horizon, response time, decision support, and confidence, are where forecasting maturity becomes tangible. They are also where the difference between stages is felt most directly by the CFO and the treasury team.



What Every Forecast Is Made Of (Hint: It's Only Three Things)

Whether you're running a simple spreadsheet or a sophisticated model, every cash forecast relies on three core ingredients:

- **Transactional and Operational Data:** Everything generated through your ledgers and banking actuals, such as customer invoices, supplier payments, banking flows, and internal postings. It reflects the real execution of your business day to day.
- **Macro Drivers:** The outside world, including interest rates, FX movements, commodity prices, and broader economic signals that shape your business environment.
- **Commercial Context:** The on-the-ground intelligence, business context, and judgment that no system can fully capture. The crucial gut-check layer.

These three sources are always there. The difference, as your forecasting matures, is how well connected, reliable, and integrated they are, and how much manual effort you're still spending trying to glue them together. Data gets interconnected, then normalized, then analyzed, then enriched; each stage building on the last until a system can drive proactive liquidity allocation.

Understanding Direct and Indirect Forecasting

Before walking through the stages, it helps to understand the two forecasting methods that sit underneath all of them. The way these methods evolve is one of the clearest indicators of forecasting maturity.

The direct method builds the forecast from actual and expected cash flows: bank transactions, invoices, receivables, payables, payroll, tax payments, and other specific movements. It is grounded in operational reality and is strongest in the near term, the 0 to 13-week horizon where transaction-level data is available and meaningful.



A Worked Example: From AR Balance to Weekly Cash Inflow

Consider a business with **\$42M in outstanding accounts receivable** and an average **DSO of 38 days**. The contractual forecast shows when cash is scheduled to arrive based on invoice due dates and expected payment terms.

But treasury also knows from historical payment behavior that the company's largest customer segment often pays **9 to 14 days later than expected**, especially near quarter-end. That means some cash that appears collectible in the near-term forecast may actually arrive later.

In this simplified example, the business expects to collect the full **\$42M** over the five-week period. However, behavioral payment patterns suggest that \$1.1M of expected collections each week will shift beyond the forecast window. However, behavioral payment patterns suggest that **\$1.1M of expected collections each week** will shift beyond the forecast window.



WEEK	AR BALANCE OUTSTANDING	CONTRACTUAL CASH INFLOW	BEHAVIORALLY ADJUSTED INFLOW	COLLECTIONS SHIFTED LATER
Week 1	\$42.0M	\$4.2M	\$3.1M	\$1.1M
Week 2	\$37.8M	\$6.8M	\$5.7M	\$1.1M
Week 3	\$31.0M	\$9.1M	\$8.0M	\$1.1M
Week 4	\$21.9M	\$12.3M	\$11.2M	\$1.1M
Week 5	\$9.6M	\$9.6M	\$8.5M	\$1.1M
TOTAL		\$42.0M	\$36.5M	\$5.5M

The timing gap is calculated by subtracting the behaviorally adjusted inflow from the contractual inflow each week:

$$(4.2-3.1) + (6.8-5.7) + (9.1-8.0) + (12.3-11.2) + (9.6-8.5) = 5.5M$$

The point is not that the company will fail to collect the cash. It is still owed the money. The issue is when. The contractual forecast shows what should happen if customers pay according to schedule.

The behaviorally adjusted forecast shows what is more likely to happen based on how customers actually pay.

Over five weeks, that timing difference creates a \$5.5M cash gap. That is a funding decision. At Stage 1, treasury plans around the contractual column and may not see the shortfall until it appears. At Stage 3, treasury uses behavioral data to anticipate the gap and manage it before it becomes a liquidity surprise.

The indirect method starts from financial planning assumptions: profit and loss projections, balance sheet movements, FP&A plans, working capital drivers, and revenue expectations. It is strongest over the medium and long term, beyond 13 weeks, where the business plan rather than individual transactions drives the liquidity outlook.

In less mature environments, treasury, FP&A, and local finance each maintain their own models and reporting cycles, built for different purposes and rarely designed to connect end-to-end.

As forecasting matures, the ideal framework compares direct and indirect views over the overlapping 1 to 3 month window, then uses the indirect method to extend the outlook. Direct forecasting gives near-term precision. Indirect forecasting gives strategic range. Together, they create a more complete picture than either could alone.



Stage 1

Fragmented Operations

Stage 1: Fragmented Operations

Technical Name: Manual and Rules-of-Thumb | Maturity Level: Ad Hoc

Stage 1 is where most treasury teams start, and where a significant portion of the market still operates today. The AFP 2024 Liquidity Survey reports that about 67 percent of treasury departments still rely primarily on spreadsheets. Kyriba's 2026 CFO Survey, covering 1,400 senior finance leaders across eight countries, found that 59 percent still lack a complete, real-time view of their cash and liquidity.

Here is what Stage 1 actually feels like from the inside: you log into six bank portals before 9 AM. You paste data into a spreadsheet someone built three years ago.



Customer PROOF

From scattered portals to one source of truth

“The transition from a manual process to the digitization of treasury with Kyriba means we can finally achieve the same-day liquidity visualization we need instead of spending our time manually populating forecasts.”

Benjamin Seal

Director of Global Treasury and Risk Management, Inotiv

[Read more](#)

You're not entirely sure how it works, but it balances, so you leave it alone. Business unit inputs arrive by email, in different formats, with different assumptions about what the word 'forecast' even means. By the time the morning position is assembled, it is already an hour stale. For multinationals operating across 15 countries with 50 or 100 bank accounts, this isn't one Stage 1 problem; it is 50 of them, each slightly different, each requiring manual intervention, none talking to each other.

Spreadsheets didn't become the dominant treasury tool by accident. They're flexible, fast to set up, and owned entirely by the team. Over time, though, many accumulate layers of legacy logic, macros and formulas built by people who may no longer be in the organization.

Teams replicate existing processes rather than challenge them, simply to avoid breaking what already functions. The forecast keeps running. The assumptions underneath it don't get reviewed.

Stage 1 is travel before GPS existed: printing directions the night before, tracing routes by hand, and pulling over to read road signs when lost. The question treasury at Stage 1 can answer, eventually, is where the money went. Not where it is going.

Are You at Stage 1?

The easiest way to spot a Stage 1 operation is the daily conversation. If you hear phrases like these, your team is right here:

- “I have to log into six bank portals before I can build the morning position.”
- “If that person is out, no one else really knows how the forecast model works.”
- “We have a 13-week forecast, but honestly we only trust the next two weeks.”
- “The CFO asked why cash was lower than expected. It took us two days to pull together the explanation.”
- “We’ve been meaning to clean up the cash categories for months but there’s never time.”

If your team spends more time building the morning cash position than using it, or if answering a basic liquidity question takes days instead of minutes, Stage 1 is your starting line. The goal isn’t to become a finance genius overnight. It’s to stop losing hours on work that a system should be doing for you.

More than half of senior finance leaders (53.4%) still rely on manual data feeds to connect their treasury systems, and 37.4% plan to fix it by simply forecasting more often, rather than fixing the infrastructure underneath.

Source: Kyriba CFO Risk Radar Survey 2026.



Stage 2

Connected Visibility

Stage 2: Connected Visibility

Technical Name: Connected Automation | Maturity Level: Emerging

Stage 2 begins with a simple premise: connect the data. Bank balances, transactions, AP schedules, AR aging, FX rates, and business unit inputs flow into a centralized environment through bank APIs and ERP connectors. Automation handles known recurring cash flows — payroll, rent, interest, tax, debt payments. The goal is a system of record: one authoritative source of truth, trusted enough that treasury stops second-guessing it and starts using it.

From that foundation, treasury can build a reliable short-term forecast for the first time, grounded in actual AP schedules, AR aging reports, and bank actuals. Business unit contributors submit through consistent templates.

The process becomes repeatable and auditable. Stage 2 handles known cash flows well. It does not yet understand behavior. If a customer usually pays twelve days late, the forecast may still show the contractual due date.

What it does unlock is meaningful: the reliable forecast window extends from two weeks to thirteen, and a liquidity question that used to take two days to answer now takes a few hours.

The direct method becomes much easier to manage. Automated bank feeds, ERP connections, and standardized cash flow categories replace the manual assembly process. A company might, for the first time, see all of its bank accounts in a single daily position across entities and currencies, without anyone spending two hours pulling it together.

Only 35% of senior finance leaders say their treasury is fully connected across banks, ERPs, and third-party data via APIs.

For the other 65%, Stage 2 isn't aspirational—it's the next concrete step.

Source: Kyriba CFO Risk Radar Survey 2026.

See connected treasury in action.



SNC

Saved up to 1 hour daily
on cash positioning.



Bolt

Moved from weekly visibility to
cash visibility every 10 minutes.



Mapletree

Connected 2,000+ bank
accounts into one platform.

The indirect method at Stage 2 is more foundational: FP&A and budget data may be imported from planning systems, but it is largely used to populate a structured template rather than drive dynamic insight. Treasury and FP&A often still maintain separate forecasts and reconcile them in meetings. That is acceptable at Stage 2. The priority is a trusted data foundation.

It is worth naming something uncomfortable about Stage 2: gaining visibility does not always reveal good news. Some organizations discover their cash position is more volatile than anyone realized, that certain entities have been carrying structural deficits, or that intercompany flows have been masking a group-level liquidity problem. Better data surfaces reality. The organizations that benefit most from Stage 2 are the ones willing to act on what it shows.



What Stage 2 Actually Delivers | SNC

SNC, a global aerospace and national security company, had reached the point where manual treasury processes could not keep pace with business complexity. With 30+ legal entities, 50+ bank accounts, operations in five countries, and four currencies, treasury needed a more connected way to see and manage cash.

After implementing Kyriba in five months, on time and under budget, SNC gained real-time cash visibility and reduced daily cash positioning work by 30 to 60 minutes every morning.

As Brian Simpson, VP Finance at SNC, said, “One of the tangible benefits we saw right away out of this implementation was real-time cash visibility, allowing us to position cash and manage debt effectively.”

SNC fits Stage 2 because the value comes from connection and visibility. Treasury is no longer spending as much time gathering bank data and assembling a cash position. The team has a reliable foundation for faster liquidity decisions.



Dory Malouf

“AI does not fix bad data foundations. It amplifies them. If your cash categories are inconsistent, your actuals are incomplete, or your business-unit inputs are unreliable, predictive models may look sophisticated while still giving the CFO the wrong answer. The foundation matters because trust is cumulative.”



A Cautionary Note What Happens When You Skip the Foundation

Here is a pattern more common than most vendors will acknowledge. A team acquires advanced forecasting capabilities, sometimes Stage 3 predictive modeling or Stage 4 AI tooling. The platform is configured. The dashboards look polished. The output is presented to the CFO with confidence. The numbers are wrong.

One pattern we see repeatedly: a team deploys a machine learning model on two years of historical cash flow data with probability weighting. It looks exactly like Stage 3. The model learned the patterns in the data perfectly. The problem is that cash flow categories had been inconsistently mapped across regions for eighteen months, a known issue that had been deprioritized. The model was trained to be confidently wrong. The output was not a better forecast. It was a worse one that looked authoritative. That is not a technology failure. That is a sequencing failure.

Better tools on top of broken inputs do not produce better answers. They produce more confident wrong answers. Stage 4 is not a shortcut around Stage 2 and Stage 3. It depends on them.

Stage 3

Intelligent Forecasting

Stage 3: Intelligent Forecasting

Technical Name: Predictive Analytics and Strategic Modeling | Maturity Level: Strategic

Stage 3 is where treasury stops reporting on the past and begins driving corporate strategy. Connected data and automated feeds gave you a trusted position at Stage 2. Stage 3 uses that foundation to ask harder questions: how accurate is our forecast, and why?

What would happen if conditions changed? How does liquidity evolve across regions and drivers over the next 12 to 18 months? And increasingly: can the system start predicting behavior rather than just recording it?

Stage 3 covers a lot of ground deliberately. Teams often move through it in a natural sequence: first building the discipline to learn from variance, then expanding into scenario modeling, then extending the horizon, and ultimately deploying AI/ML-powered prediction. A team in early Stage 3 and a team in advanced Stage 3 are doing fundamentally different things, but they're on the same journey.

Learning from Variance

The defining Stage 3 capability is structured, ongoing variance analysis; not the retrospective 'why were we off?' conversation that happens after a CFO escalation, but a systematic comparison where treasury tracks Mean Absolute Percentage Error (MAPE) by category and entity. "PwC's 2025 Global Treasury Survey found that 76 percent of treasury teams cite poor data quality as a forecasting challenge even after implementing connectivity tools.

Without a feedback mechanism, there is no way to know whether AP is consistently off by 15 percent, whether a specific business unit systematically underestimates collections, or whether a seasonal pattern has shifted since the model was last calibrated.

At Stage 3, the system surfaces actuals versus forecast automatically: versioned, time-stamped, and broken down by category. Treasury identifies where variance is widest, shares accuracy data back to business unit contributors, and the forecast improves over time because the process is designed to learn. Performance KPIs, such as forecast accuracy by category, variance trends, and version comparisons, become part of the regular reporting rhythm.

32.5% of senior finance leaders identify forecast accuracy and data reliability as a top concern, yet fewer than 1 in 4 (23.3%) list forecasting accuracy as a named treasury priority.

The gap between concern and action is where Stage 3 lives.

Source: Kyriba CFO Risk Radar Survey 2026.

Scenario Testing and Long-term Planning

The question treasury begins to answer at Stage 3 is no longer just ‘what will happen’ but ‘what if.’ Scenario modeling allows treasury to build, compare, and stress-test multiple liquidity pictures simultaneously. Sensitivity analysis reveals which assumptions have the most impact on the liquidity runway. Treasury can walk into the CFO’s office with a range of outcomes, the drivers behind each, and a clear view of where the business is most exposed.

Stage 3 also extends the forecast horizon from weeks to months and eventually to a year or more. Bottom-up submissions from regional entities, rolling multi-month forecasts, debt maturity planning, FX exposure modeling, and capex sequencing all become part of a connected liquidity picture. The indirect method becomes dynamic: static budget imports are replaced by driver-based models where FP&A assumptions and revenue plans translate directly into projected cash impact. When revenue is tracking above plan, the forecast immediately reflects the working capital implications — long before it hits the bank statements.

The most powerful new input at Stage 3 is the operational cash signal: the behavioral data that lives inside the AP and AR workflow itself. Invoice status, approval timing, payment terms, early-payment discounts, and supplier-specific payment history are leading indicators of when cash will actually move — not the contractual date, but the behavioral one. When that workflow-level intelligence feeds the forecast directly, treasury stops inferring behavior from bank actuals after the fact and starts anticipating it at the source. This is the input layer that turns a forecast that records behavior into one that predicts it.



Cooke Aquaculture's treasury team needed to move beyond fragmented visibility and manual forecasting processes toward a more mature, data-driven liquidity planning model. Before advancing its forecasting capabilities, the team first strengthened the foundation: improving cash visibility, standardizing processes, and creating a more reliable view of liquidity across the business.

With Kyriba, Cooke improved real-time cash visibility from 65 percent to 100 percent and increased forecast accuracy to approximately 90 percent. The team also raised its treasury maturity score from 1.93 to 3.57 on a 4-point scale, while reducing annual manual workloads from 8,320 hours to 1,414 hours, an 83 percent reduction.

Cooke fits Stage 3 because forecasting becomes a learning process, not just a reporting exercise. Better visibility gives treasury a trusted baseline. Variance analysis helps the team understand where forecasts are accurate and where assumptions need adjustment. Over time, the process improves, giving treasury more capacity to support procurement, capital planning, operations, and strategic liquidity decisions.

Stage 3 is the GPS that does more than show you where you are. It has learned from your previous journeys: which routes you prefer, where delays tend to happen, and how long your trips usually take. Now it is planning ahead. Three different route options for Friday's commute. A probability estimate for when you'll actually arrive based on historical traffic patterns. For treasury, that is the difference between a team that produces a forecast and a team that continuously improves one.



What Stage 3 Sounds Like

If your team is operating at Stage 3, the conversation changes in ways you can hear:

- “AR was off by 12 percent last month. We know it’s the segment that tends to slip at quarter-end. We’ve adjusted the model.”
- “We ran three scenarios for the CFO. The base case works, but the downside scenario shows a funding gap in Q3 if collections slow.”
- “Our indirect forecast now covers 18 months. When the acquisition timeline shifted, we had the liquidity model updated within the day.”
- “FP&A’s revenue plan is tracking above forecast, but collections haven’t moved yet. We surfaced it before the month closed.”

Diagnostic for Stage 3

If your team produces a forecast every week but has no systematic view of how accurate last week’s was, or if scenario testing still means building a separate spreadsheet by hand each time the CFO asks a ‘what if’ question, Stage 3 is your next move.

The real question at Stage 3 is whether treasury is helping the business make decisions, or just reporting on them.



Stage 4

AI Guided Liquidity Decisions

Stage 4: AI-Guided Liquidity Decisioning

Technical Name: Agentic AI & Governance | Maturity Level: Leader

Stage 4 is the future, but it's not science fiction. At this stage, the system reasons across all available information: monitoring continuously, detecting when something looks off, identifying root causes, modeling alternatives, checking recommendations against policy, and handing treasury a ready-to-act analysis. This is agentic AI applied to treasury.



**Edouard
Gabreau**

“The practical value of AI in treasury is not replacing judgment. It is reducing the time between signal and decision. If the system can surface a variance, explain the driver, check the policy, and suggest next steps, treasury can spend more time deciding and less time searching.”

Want to see this in action? See how HelloFresh uses Kyriba TAI to reconcile hard-to-find payments in seconds.

[Read more](#)

The shift is worth naming carefully. At Stage 4, you go from running every step of the process to governing the process.

92% of senior finance leaders are already using AI in at least some processes. More than half (53%) list it as a top treasury priority.

The question is no longer whether to adopt AI—it's whether the foundation is ready to use it well.

Source: Kyriba CFO Risk Radar Survey 2026.

That sounds like less work. Governance is actually harder than execution; deciding what the system can recommend, what it can initiate autonomously, and what always requires human sign-off demands more judgment than building a spreadsheet did. Stage 4 doesn't simplify treasury. It elevates it.



What Stage 4 Looks Like in Practice

Here is what the capability actually produces:

“Your 30-day liquidity position is projected to be \$12 million below last week’s forecast. The primary driver is delayed receivables from three customers whose payment behavior shifted over the past ten days. At the current trajectory, the minimum liquidity threshold may be reached in 18 days. Recommended actions: draw \$8 million on the revolver before Thursday’s rate reset, initiate collections outreach to the three customers, and defer the planned EMEA capital expenditure by 30 days. All recommendations are within approved treasury policy limits.”

No one spent the night building that analysis. The system identified the variance, traced it to specific drivers, modeled alternatives, checked policy limits, and prepared a recommendation for review. Treasury’s role is not to build the analysis; it is to evaluate it, challenge it where the model may be missing commercial context, and decide.

The AI intelligence at Stage 4 also goes beyond history-based pattern recognition. The system distinguishes between what has happened before and what is beginning to happen now, detecting behavioral shifts in payment patterns, flagging anomalies against peer benchmarks, and incorporating external signals that affect liquidity long before they appear in the ERP: commodity price indices, shipping lead times, interest rate curves, or weather patterns that shift customer buying behavior and collections timing.



What Remains Human at Stage 4

Stage 4 doesn't make treasury less important. It makes judgment more visible and more consequential.



Edouard
Gabreau

“The more intelligent the system becomes, the more important human governance becomes. Treasury still owns policy, judgment, risk appetite, bank relationships, and exception handling. AI can prepare the recommendation. Treasury decides whether the recommendation is right for the business.”

Policy, risk appetite, bank relationships, counterparty judgment, exception handling, and major decisions stay with people. Treasury leaders decide exactly what the system can recommend, what it can initiate autonomously, and what always requires human sign-off. Autonomy without governance is not innovation. It is risk.

Analysts move into roles as interpreters, challengers, and decision partners, injecting the commercial and operational context that models cannot encode. Rather than simply validating what the system produces, they improve what the system is trying to decide. Think of Stage 4 less like a robot treasurer and more like a highly capable co-pilot: watching everything constantly, flagging risk early, always having options ready. But treasury is still flying the plane.

There is an important shift embedded in a recommendation like the one above. Seeing the cash gap is insight; closing it requires action like outreach to specific customers, a change to payment timing, a deferred outlay. Historically those levers have lived outside treasury, in the AP and AR workflows treasury could observe but not direct. Stage 4 only delivers its full value when insight and actionability are joined: the forecast identifies what should change, and the operational workflow has the muscle to execute it. A recommendation no one can act on is just a better-looking report. The organizations that pull ahead are the ones that connect the decision to the workflow that carries it out.

Good treasury professionals develop something no system can replicate: judgment built from experience. They understand the nuances of their markets, the business's cash cycles, the counterparties that matter, and the decisions that carry long-term consequences. At Stage 4, that judgment isn't replaced. It is amplified.



A Practical Maturity Path for Treasury Teams

The path forward will not look identical for every organization. Some companies have excellent bank connectivity but weak AR forecasting. Others have strong FP&A planning but poor treasury visibility. Some track variance but have never run a formal scenario model. Some have predictive models but no governance structure for acting on what those models recommend. The framework is not meant to label teams as mature or immature. It is meant to help treasury leaders identify the next best move.

The simplest way to start is to ask where the friction is, and where you want to be.

CURRENT CAPABILITY PATTERN	LIKELY STAGE	FORECAST HORIZON TODAY	RECOMMENDED FOCUS
Cash position built manually from multiple bank portals	Stage 1	0–2 weeks	Establish bank connectivity, standardize cash categories, centralize the daily cash position
Cash position centralized but forecast built in spreadsheets	Stage 1 / Stage 2 starting	2–4 weeks	Automate ERP feeds, standardize AP and AR data, move forecast into a connected environment
Connected visibility achieved; 13-week forecast in place	Stage 2 complete	Up to 13 weeks	Introduce variance tracking; begin measuring forecast accuracy by cash flow category
Data connected and automated; no systematic variance analysis	Stage 2 / Stage 3 starting	13 weeks	Build actuals-vs-forecast comparisons; track accuracy by category; establish performance KPIs
Variance tracking exists; no scenario modeling or long-term view	Early Stage 3	13–26 weeks	Introduce scenario versions; build sensitivity analysis; extend horizon to 6 months
Direct cash forecast and FP&A plan exist but are maintained separately	Mid Stage 3	3–6 months	Reconcile operational forecasts with planning assumptions; extend indirect horizon to 12–18 months
Scenario modeling and long-term planning in place; forecast is still manually built	Advanced Stage 3	6–18 months	Introduce behavioral AP/AR forecasting; connect external data signals; prepare AI/ML foundation
Predictive models are running; recommendations still built manually	Advanced Stage 3 / Stage 4 starting	12–18 months	Build the governance model; define what the system can recommend, initiate, and escalate
System detects issues, explains drivers, and recommends actions	Stage 4	12–24 months rolling	Strengthen oversight; expand external data signals; build continuous learning loops

Treasury as Strategic Co-Pilot

The future of forecasting isn't about getting a prettier report. It's about securing a completely different role for treasury. At Stage 1, the team is assembling the picture, but by the time it's finished, the moment to act has often passed. At Stage 2, you get control of your data, cutting response time from days to hours. At Stage 3, the real analytical work begins: variance discipline, scenario building, and AI-powered prediction, moving the team from reporting on the past to actively shaping decisions about the future. At Stage 4, treasury gains an intelligent co-pilot that monitors continuously and has a recommendation ready before the question is asked.

For analysts, this means less time in portals and fewer late nights reconciling numbers that should have matched in the first place; more time thinking, challenging assumptions, and helping the business decide what to do next. For CFOs, it means fewer surprises, faster access to trusted answers, and the ability to make strategic decisions in hours rather than days.

"The implementation of Kyriba demonstrates that collaboration between technology, people, and processes is the true engine of transformation. We are moving from an operational treasury to a strategic treasury—solid and future-ready."

Asier Pe-Mentxaka Fernández
Corporate Finance Lead, Exolum

[Read more](#)

Only 38% of senior finance leaders describe treasury as a strategic growth driver. Nearly 1 in 3 (30.7%) still see it as a support function or back-office cost center.

Source: Kyriba CFO Risk Radar Survey 2026.

HelloFresh operates across more than 16 markets and generates nearly \$8 billion in annual revenue, while its group treasury function remains lean. That scale requires technology that can do more than report data. It needs to help treasury answer questions, resolve issues, and move faster.

Kyriba's Trusted Agentic AI, TAI, gives HelloFresh a practical example of AI as a treasury co-pilot. Felix Luber describes a payment reconciliation issue that every treasurer recognizes: a payment that cannot be matched easily and takes hours to investigate. With TAI, treasury can ask the system to find and reconcile the payment, and the answer comes back within seconds. He also describes using TAI to answer audit and documentation questions that once required manual searches through technical files.

HelloFresh is actively using AI to help treasury perform work, answer questions, and remove operational bottlenecks.

The treasury team shifts from creating and searching for data to interpreting results, governing actions, and making decisions.

The Liquidity Intelligence Gap Is Not a Technology Gap

The organizations that close the liquidity intelligence gap in the next 24 months will have a structural advantage in capital efficiency that latecomers cannot simply buy back. This is not primarily a software procurement decision. It is a governance decision — whether your organization is willing to build the foundation before reaching for the capability, and whether leadership is prepared to act on what better intelligence actually shows.

The better measure of forecasting maturity is behavioral: has treasury stopped explaining the past and started shaping the future? Has the CFO started acting on scenarios rather than waiting for certainty? Has the forecast changed a decision in the last 90 days that would have gone differently without it? That is the destination: not a single technology purchase, but the continuous journey from fragmented operations to autonomous liquidity leadership. Every stage builds the foundation for the next.

Where Does Your Team Stand Today?

Before your next planning cycle, before the next technology evaluation or CFO conversation about forecast accuracy, schedule time with our Business Value Advisory team to better quantify, benchmark, and accelerate the value that maturing your forecasting capabilities would provide to your organization.

The Business Value Assessment is complimentary.
[Request a Performance Maturity Assessment.](#)

