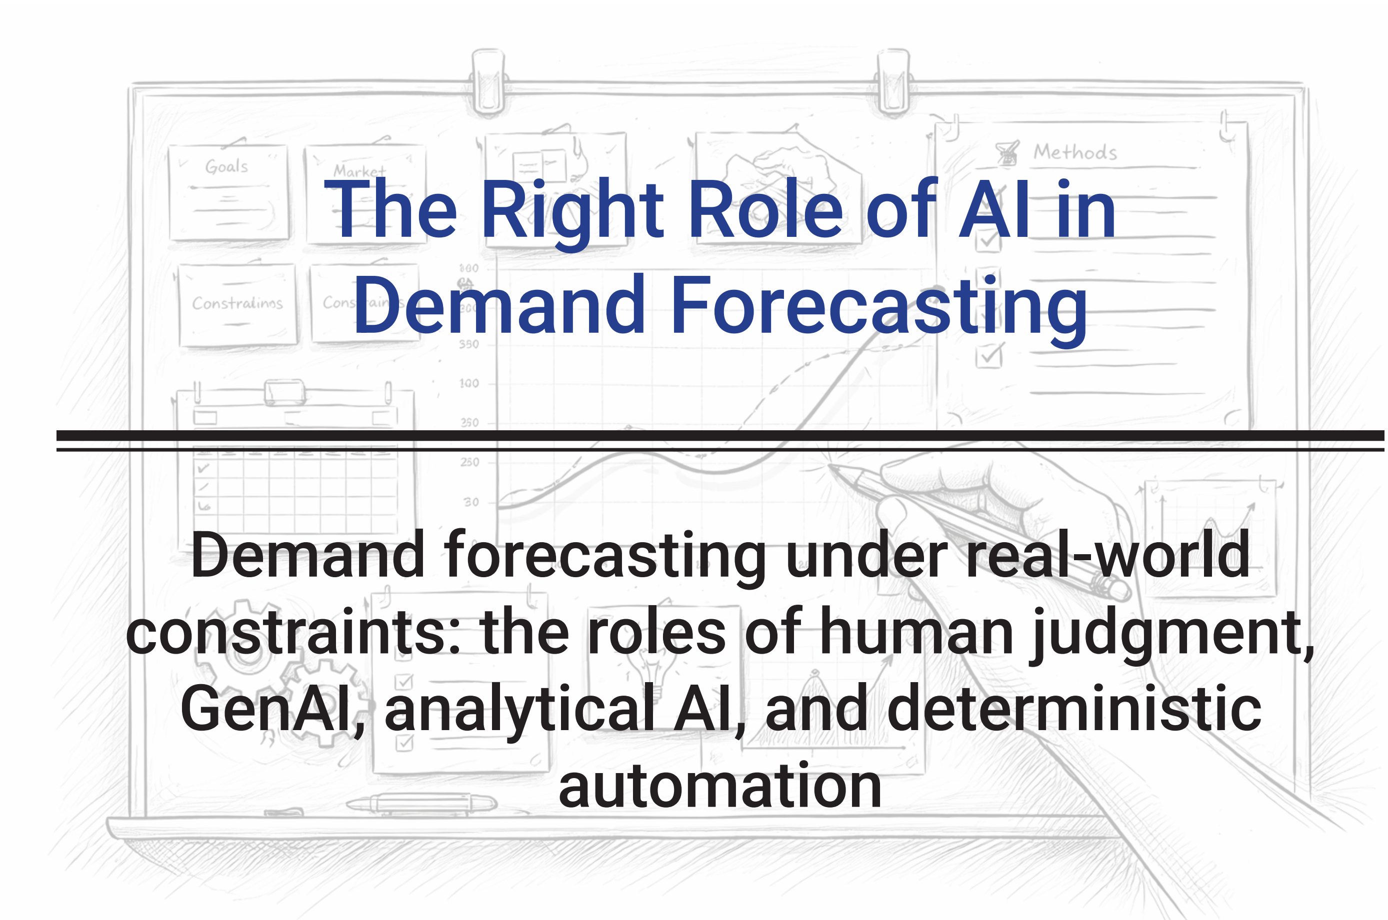

The Right Role of AI in Demand Forecasting

Demand forecasting under real-world constraints: the roles of human judgment, GenAI, analytical AI, and deterministic automation

Executive summary

Demand forecasting is often discussed as a modeling problem, but in practice it is a broader decision process. Before any forecast is produced, the business must decide what should be forecast, why it matters, what level of detail is worth the effort, and what actions will depend on the result. After that, it still has to gather relevant inputs, choose an approach, select a method, execute the forecast, and monitor whether the result remains fit for use over time. In other words, demand forecasting is not just about prediction. It is about assigning the right kind of work to the right kind of capability.

This paper analyzes that process step by step and distinguishes four forms of support: human judgment, GenAI, analytical AI and statistical methods, and deterministic automation. The core argument is simple. Human experts should continue to own decisions that are specific to the business, especially where priorities, trade-offs, sufficiency, and action must be judged in context. GenAI is most useful as a support tool: helping teams clarify questions, compare options, identify obvious gaps, explain results, and reduce routine technical effort in areas such as data preparation and method implementation. Analytical AI, machine learning, and statistical methods are most useful where the task is actual prediction, pattern detection, or quantitative comparison. Deterministic automation remains essential wherever the work is rule-based, repeatable, and should be consistent.

A clear pattern emerges across the process. Early steps are dominated by framing and choice, which makes human judgment central. Middle steps combine human decisions with stronger roles for GenAI, analytical methods, and automation, especially when preparing data and implementing candidate methods. Later steps shift toward execution, monitoring, and control, where analytical methods and deterministic automation become much more important, but still do not remove the need for human review. This allocation matters because many organizations make the same mistake in different forms: they ask AI to compensate for an unclear question, weak process discipline, or undeclared business assumptions.

The practical conclusion is not that GenAI should be avoided, nor that it should be inserted everywhere. It should be used where it broadens thinking, reduces omission risk, and removes mechanical effort without taking ownership away from the right actor. It should not be used as the owner of the business question, as a substitute for deterministic controls, as a replacement for forecasting methods, or as the final judge of correctness, completeness, or business readiness. In demand forecasting, the strongest results come not from maximizing AI usage, but from keeping judgment human, using analytics where prediction is required, and applying automation where consistency is enough.

What is this white paper about?

This analysis is not an exercise in “where can we use AI”. It is an exercise in understanding the real work inside a business process, and then deciding which kind of capability fits each part of that work.

The goal is to separate five different things that are often mixed together:

- **Decision work**: choosing goals, trade-offs, methods, thresholds, and actions
- **Knowledge work**: gathering practices, comparing options, identifying gaps, explaining choices
- **Prediction work**: estimating future outcomes from data
- **Execution work**: collecting, transforming, organizing, documenting, and presenting information
- **Control work**: checking completeness, consistency, quality, and readiness

This distinction matters because different kinds of work need different kinds of support.

Core principle

The human expert remains responsible for decisions that are specific to the organization and its context.

This includes decisions such as:

- why the work is being done
- what outcome matters
- which method is appropriate here
- whether the work is complete enough
- what conclusions should be drawn
- what actions should follow

These are not generic decisions. They depend on local constraints, business priorities, risk tolerance, and the meaning of success in this specific situation.

Role of GenAI

GenAI can support the process in two ways:

- Cognitive support: finding practices, comparing options, critiquing reasoning, summarizing, explaining
- Implementation support: generating code, scripts, transformations, tests, reports, and other technical artifacts needed to execute the chosen approach

As cognitive support, it helps teams:

- find known practices, common approaches, and expert recommendations
- surface alternative options that might otherwise be missed
- identify obvious gaps, weak assumptions, and incomplete reasoning
- summarize, structure, rewrite, and explain information
- reduce routine work around collection, transformation, preparation, and presentation of information

GenAI is useful for improving the quality and speed of thinking-support work. It helps reduce the risk of narrow, careless, or poorly informed work.

As implementation support, tools such as code-oriented GenAI assistants can help create the technical artifacts needed for forecasting work, including:

- data preparation scripts

- statistical model code
- backtesting routines
- validation checks
- reporting logic
- experiment scaffolding
- pipeline components

In both cases, GenAI supports the work but does not own the decision. It does not determine what is right for this organization. Its outputs are probabilistic, pattern-based, and dependent on prompt quality, context quality, and review.

Role of other AI methods

Other AI methods are used where the task is primarily analytical or predictive.

This includes work such as:

- forecasting demand
- detecting patterns and anomalies
- classification
- optimization
- segmentation
- statistical modeling
- time-series analysis

These methods are useful when the main challenge is not explanation or drafting, but estimation, pattern recognition, or quantitative inference from data.

Role of deterministic automation

Not all useful automation is AI. Some parts of the process are best handled through deterministic automation, such as:

- rules
- workflows
- validation checks
- BI logic
- SQL transformations
- ETL pipelines
- scheduled reports
- threshold alerts

This category is important because it is often more reliable, more transparent, and easier to govern than probabilistic systems. If a task can be handled well by deterministic automation, there should be a strong reason to replace it with AI.

Human error, AI uncertainty, and reliability

This analysis also recognizes that each actor has different strengths and weaknesses.

Humans bring contextual judgment, accountability, and the ability to decide what matters here. They also bring inconsistency, bias, fatigue, and avoidable mistakes.

GenAI brings breadth of knowledge, semantic flexibility, and speed in language-based tasks. It also brings uncertainty, hallucination risk, and probabilistic output.

- **Deterministic automation** brings consistency and repeatability, but only within clearly defined rules.
- **Predictive AI and ML** can detect patterns and produce forecasts beyond manual analysis, but they depend on data quality, modeling assumptions, and proper interpretation.

The point of this analysis is not to choose one over the others. It is to assign each kind of work to the form of support that fits it best.

What we evaluate at each step

For each step in the process, the analysis asks:

- What is the purpose of this step?
- What kind of work is happening here?
- What requires human judgment?
- Where can GenAI help?
- Where are other AI methods more appropriate?
- Where is deterministic automation enough?
- What are the main risks if this step is done poorly?
- Who should own the final decision at this step?

Decision rule behind the analysis

The guiding rule is simple:

- Use **human judgment** where the choice is specific to this business
- Use **GenAI** where broad knowledge, critique, explanation, or routine language work is needed
- Use **predictive or analytical AI** where the core task is estimation, pattern recognition, or optimization
- Use **deterministic automation** where the work is rule-based and repeatable

The purpose is not to maximize AI usage. The purpose is to improve the quality, speed, and reliability of the process without losing control of decisions that must remain human.

Reference process used in this analysis

This analysis uses the student-oriented demand forecasting process described in *Supply Chain Management: An Integrated Approach*, section **“9.2 Demand Forecasting Process,”** published on Pressbooks. We selected this source deliberately because it presents demand forecasting as a standard textbook process rather than as a vendor-specific workflow or a company-specific operating model. It is the kind of structured process many students and practitioners are taught first when learning the basics of forecasting in logistics and supply chain work.

The process steps:

What to forecast: Define the forecasting target and horizon: long-term, medium-term, or short-term. The source ties this directly to business purpose and planning horizon.

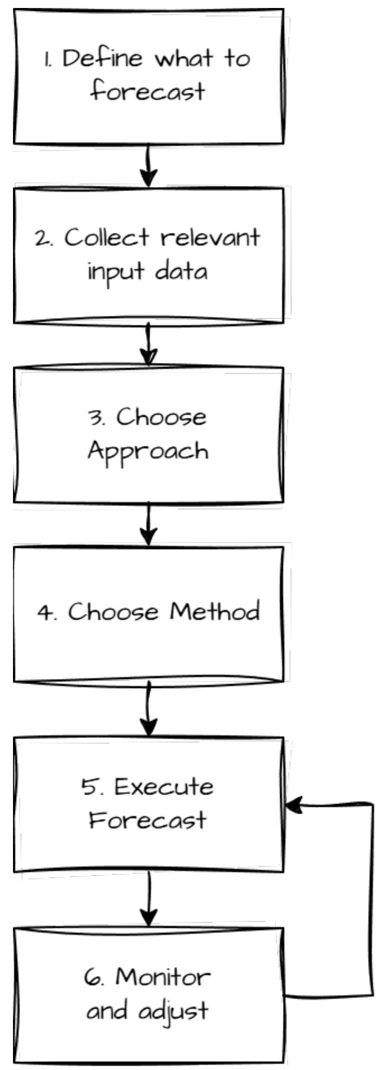

Data collection: Gather the inputs needed for forecasting, such as historical demand, economic indicators, demographic or trend data, and expert knowledge.

Forecasting approach: Choose the broad approach: qualitative (judgmental) or quantitative. The text explains that this depends mainly on data availability and the nature of the forecast.

Method selection: Select the specific method within that approach. For quantitative forecasting, the text says the method is typically chosen based on fit with historical demand patterns; for qualitative forecasting, it depends on decision type, expertise, and horizon.

Forecast execution: Run the forecast using the selected method. The same section also notes that forecast accuracy should be measured by comparing predictions with actual outcomes.

Monitoring: Continuously monitor the forecast and adjust the method or parameters as conditions change. The text explicitly says a model that was once suitable may become inadequate over time.

Actual processes in real organizations may differ in detail. Companies may distribute responsibility differently, combine or split some steps, automate some parts more heavily, or apply different levels of rigor depending on the business context, forecasting horizon, and operational stakes. However, those variations usually still map to the same underlying structure: define the forecasting need, gather inputs, choose the logic of approach, select a method, produce the forecast, and review whether it remains fit for use. That is why this textbook model is a useful foundation for a general analysis.

This document therefore does not assume that every organization performs demand forecasting in exactly the same way. It assumes something more practical: while implementation details vary, the core logic of the process is stable enough to support a step-by-step analysis of where human judgment, GenAI, other analytical methods, and deterministic automation fit best.

Process Steps

Step 1: Define what to forecast

What happens in this step

The business defines the forecasting target and the planning horizon, based on the specific needs and objectives of the business, including whether the horizon is long-term, medium-term, or short-term. The forecasting method should match the purpose and horizon of the forecast.

Why this step exists

This step exists to prevent the organization from building a technically sound forecast for the wrong question. Before collecting data or selecting a method, the business needs to be clear about what is being forecast, for what decision, and over what time period.

Type of work

Primarily:

- Decision work
- Knowledge work

Secondarily:

- Control work

This is not yet prediction work in the technical sense. At this point, the main task is defining intent, scope, and meaning.

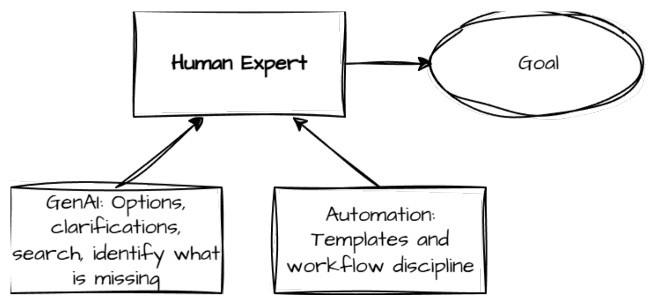

What requires human judgment

This step is strongly human-led.

Human judgment is needed to decide:

- what business question actually matters,
- what object should be forecasted, for example total demand, product family demand, SKU-level demand, region-level demand, or channel-level demand,
- what horizon matters, short-term, medium-term, or long-term,
- what level of granularity is worth the cost and complexity,
- what the forecast will be used for, inventory, staffing, budgeting, pricing, capacity planning, or strategic bets,
- what “good enough” means in this context.

This is exactly the kind of work that is specific to the organization. A model cannot decide this correctly on its own, because the right target depends on local priorities, operating constraints, and the consequences of being wrong.

Where GenAI can help

GenAI can be useful here, but as support, not as owner.

Useful roles for GenAI in this step include:

- helping articulate the forecasting question clearly,
- turning a vague request into several possible forecast definitions,
- showing how other organizations commonly frame similar forecasting problems,
- comparing different forecast scopes and horizons,

- surfacing obvious missing considerations, for example whether the forecast is meant for operational replenishment or strategic planning,
- drafting a structured problem statement,
- translating between business language and analytics language.

This fits our agreed model well: GenAI helps teams not miss known practices, typical options, and obvious gaps. It improves the quality of framing, but it does not determine what is right for this business.

Where other AI methods can help

Very little at this stage, directly.

Classical AI, ML, or statistical methods are usually not yet the main tool here, because the forecasting problem has not been defined well enough. They may help only indirectly, for example:

- by showing whether a proposed level of granularity is realistic given available data,
- by testing whether certain targets are even forecastable with acceptable accuracy,
- by helping compare alternative target definitions later.

But in the first step, predictive AI is not the center of gravity. The main issue is problem definition, not computation.

Where deterministic automation is enough

Some support tasks can be handled well through deterministic automation, for example:

- standard intake forms for forecast requests,
- templates that force definition of target, horizon, use case, and decision purpose,
- workflow checks that prevent moving forward without required fields,
- rule-based classification of forecast horizon or planning category,
- standard documentation and approval flow.

This is a good example of where not every useful improvement needs AI. If the organization repeatedly forgets to define scope, horizon, or business purpose, a structured template may solve more than a chatbot.

Risks and failure modes

This step carries major downstream risk.

What can go wrong:

- the business forecasts the wrong thing,
- the horizon does not match the decision being supported,
- the forecast is too aggregated to be actionable,
- the forecast is too granular to be reliable or economical,
- the team confuses interest with decision value and forecasts what is easy instead of what matters,
- GenAI produces a polished but weak framing that sounds reasonable and hides ambiguity,
- the organization starts discussing methods before agreeing on the actual question.

This is a classic false-confidence step. If the framing is wrong, everything downstream may still look rigorous while being aimed at the wrong target.

Allocation summary

Best fit: Human-led, with GenAI support and deterministic structure.

Step 2: Collect relevant input data

What happens in this step

The business gathers the information needed to support the forecast. This step includes historical sales data, economic indicators, demographic information, trends, and expert opinions. The type of data needed depends on the forecasting approach and the business context.

Why this step exists

This step exists because a forecast cannot be better than the inputs and assumptions that support it. Once the forecasting target and horizon are defined, the organization needs to assemble the data that may actually explain or predict demand for that target. It is a distinct step before choosing the approach and method, which is sensible: without understanding what data is available and usable, method choice is premature.

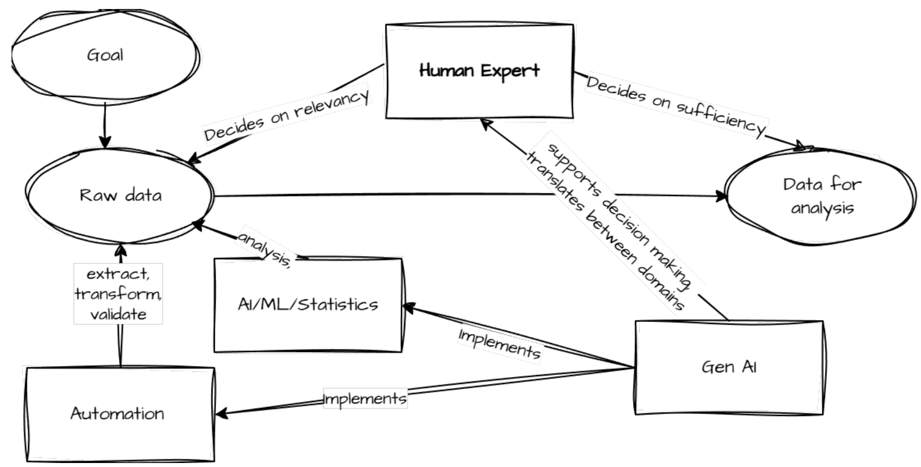

Type of work

Primarily:
- Execution work
- Knowledge work

Secondarily:
- Control work

This step is mostly about assembling, understanding, and validating candidate inputs. It is not yet the forecast itself.

What requires human judgment

Human judgment is still essential here.

It is needed to decide:
- which data sources are actually relevant to the forecasting target,
- which inputs reflect real business drivers and which are noise,
- whether expert judgment should supplement missing or unreliable data,
- whether available data is consistent enough for the intended use,
- how much effort is justified to collect additional inputs,
- whether the organization is collecting data because it is useful or merely because it is available.

This is another place where “what is right for us” matters. The business has to judge not only data availability, but business meaning. A dataset may be clean and complete and still be useless for the decision.

Where GenAI can help

Cognitive support

GenAI can be quite useful in this step.

Useful roles include:
- helping inventory possible input sources,
- suggesting common categories of data used in similar forecasting situations,
- explaining what different data fields may mean in business terms,

- comparing internal data needs against known practices,
- identifying likely gaps in the current data picture,
- summarizing subject-matter expert input,
- extracting structured notes from interviews with planners, sales, or operations,
- drafting a data requirements list,
- translating between business teams and analytics teams.

This is a good fit for GenAI because much of the work involves structuring messy information, surfacing known patterns, and reducing omission risk.

Implementation support

Code-oriented GenAI tools can also accelerate the technical work required to prepare data for forecasting. This may include:

- writing SQL queries
- building data extraction and transformation scripts
- preparing feature engineering code
- generating validation checks
- creating exploratory notebooks
- drafting reproducible preprocessing pipelines
- producing data quality checks and diagnostic summaries

This can reduce routine technical effort and speed up iteration. However, GenAI should not be trusted to define data relevance, business meaning, or sufficiency on its own. A script that runs successfully may still prepare the wrong inputs, apply the wrong definitions, or hide weak assumptions. Human review is still required to confirm that the prepared data matches the forecasting purpose.

Where other AI methods can help

Other AI and analytical methods can start to matter here, but mostly in support of data assessment rather than forecasting itself.

Useful roles may include:

- anomaly detection on historical demand data,
- clustering or segmentation to identify distinct demand patterns,
- statistical checks for missingness, volatility, seasonality, or structural breaks,
- feature relevance testing,
- exploratory analysis to see whether candidate inputs appear predictive.

This is still not the main predictive step, but analytical methods can already help separate promising inputs from weak ones.

Where deterministic automation is enough

A lot of this step is a strong candidate for deterministic automation.

Examples include:

- data extraction from systems of record,
- ETL pipelines,
- scheduled imports,
- schema validation,
- completeness checks,
- duplicate detection,
- unit normalization,
- standard joins and transformations,
- exception reporting,

- standard data quality dashboards.

This matters a lot. Much of "data collection" is not a GenAI problem at all. It is disciplined data engineering, workflow, and validation.

Risks and failure modes

Common problems in this step include:
- collecting data that is available but not relevant,
- using inconsistent definitions across systems,
- missing key explanatory factors,
- relying on expert opinion without making assumptions explicit,
- pulling historical data that does not match the current business reality,
- underestimating data quality problems,
- letting GenAI summarize flawed source material in a polished way that hides weakness,
- confusing data volume with data usefulness.

There is also a classic governance risk here: teams often move forward because they have some data, not because they have the right data.

Allocation summary

Best fit: Deterministic automation for collection and validation, human judgment for relevance and sufficiency, with GenAI support for structuring, gap-finding, and cross-functional translation.

Step 3: Choose the forecasting approach

What happens in this step

The business chooses the broad forecasting approach. This step is choosing between two primary approaches: qualitative and quantitative. Qualitative forecasting as useful when historical data is unavailable or not relevant, and quantitative forecasting as appropriate when past patterns are expected to recur, especially for short- and medium-term forecasts.

Why this step exists

This step exists because before selecting a specific method, the organization needs to decide what kind of logic will govern the forecast. This is the point where the business decides whether the problem is better treated as a judgment-heavy question, a data-driven pattern question, or some combination of both.

Type of work

Primarily:
- Decision work
- Knowledge work

Secondarily:
- Control work

This is still not the stage where the forecast is actually produced. It is about selecting the right family of approach for the nature of the problem.

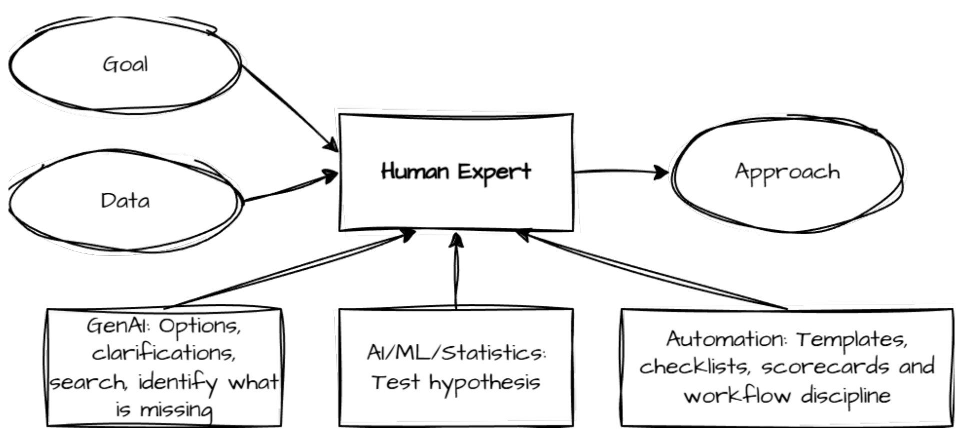

What requires human judgment

This step is strongly human-led.

Human judgment is needed to decide:

- whether historical data is relevant enough to justify a quantitative approach,
- whether the business is facing a situation where expert judgment matters more than pattern repetition,
- whether the horizon and business context make quantitative forecasting realistic,
- whether a hybrid approach is more appropriate than a pure one,
- whether the organization has the data quality, stability, and process discipline needed to support a quantitative approach,
- whether the decision is too high-stakes to rely on one style of reasoning alone.

This is another step where "what fits us" matters more than "how others do that." A business can easily choose a quantitative approach because it feels more rigorous, even when the data is thin, distorted, or no longer representative. It can also overuse qualitative judgment where the problem is actually structured enough for more disciplined modeling.

Where GenAI can help

GenAI can be useful here as a thinking aid and critic.

Useful roles include:

- explaining the practical difference between qualitative and quantitative approaches,
- showing common criteria used to choose between them,
- summarizing typical pros, cons, and failure modes,
- helping the team compare several candidate approaches,
- surfacing obvious mismatches, for example trying to force quantitative forecasting onto a situation with weak or irrelevant historical data,
- helping draft the rationale for why one approach was selected,
- translating statistical or forecasting concepts into plain business language.

This is a good fit for GenAI because the work is still heavily language- and reasoning-based. It helps the team not miss standard options and obvious weaknesses in its own thinking.

Where other AI methods can help

Other AI and analytical methods can help here, but mainly as evidence for the choice, not as the choice itself.

Useful roles may include:

- testing whether historical data contains stable enough patterns to support quantitative forecasting,
- evaluating signal strength in available data,
- identifying whether meaningful segmentation exists,
- assessing whether demand behaves in ways that are likely to be forecastable,
- comparing baseline model performance across simple candidate approaches.

In other words, predictive and analytical methods can help answer, "Is a quantitative path viable here?" But they do not replace the business decision about whether that path is appropriate for the situation.

Where deterministic automation is enough

Some support can be handled well through deterministic automation, for example:

- decision templates that require the team to document horizon, data availability, stability, and use case,
- checklists that force explicit consideration of qualitative vs quantitative criteria,
- workflow gates that prevent method selection before approach rationale is documented,
- standard scorecards for data readiness,
- rule-based routing of simple cases into predefined forecasting paths.

This is important because some discipline in this step comes not from AI, but from forcing the organization to make its reasoning explicit.

Risks and failure modes

Common problems in this step include:

- choosing quantitative forec1asting because it looks more objective, even when the data does not justify it,
- choosing qualitative forecasting because it feels safer, even when the problem is structured enough for stronger analytical treatment,
- treating the choice as binary when a hybrid setup would work better,
- confusing data availability with data relevance,
- letting GenAI produce a polished comparison that sounds persuasive but hides shallow reasoning,
- choosing an approach based on internal capability or politics rather than on the nature of the problem,
- skipping explicit rationale and moving straight to favorite methods.

This step is especially vulnerable to false rigor. Teams often adopt the approach that sounds more professional, rather than the one that actually fits the decision and the evidence.

Allocation summary

Best fit: Human-led, with GenAI support for comparison and critique, analytical methods for viability testing, and deterministic structure to make the rationale explicit.

Step 4: Select the forecasting method

What happens in this step

The business selects the specific forecasting method within the chosen approach.

Why this step exists

This step exists because “qualitative” and “quantitative” are still too broad to execute. The organization now has to decide which specific method is appropriate for the target, the horizon, the data, and the decision context. This is where the process moves from general logic to concrete technique.

Type of work

Primarily:

- Decision work
- Knowledge work

Secondarily:

- Control work
- Prediction work

This is the point where analytical and statistical considerations start to become more central, but the main task is still selecting rather than forecasting.

What requires human judgment

Human judgment is essential here.

It is needed to decide:

- which specific method fits the shape and stability of the demand pattern,
- whether a simple method is sufficient or a more complex one is justified,
- whether interpretability matters more than marginal accuracy gains,
- whether the business can actually support the chosen method operationally,
- whether the method matches the quality and volume of available data,
- whether the method fits the consequences of forecast error,
- whether the team is choosing a method because it is appropriate, or because it is fashionable or familiar.

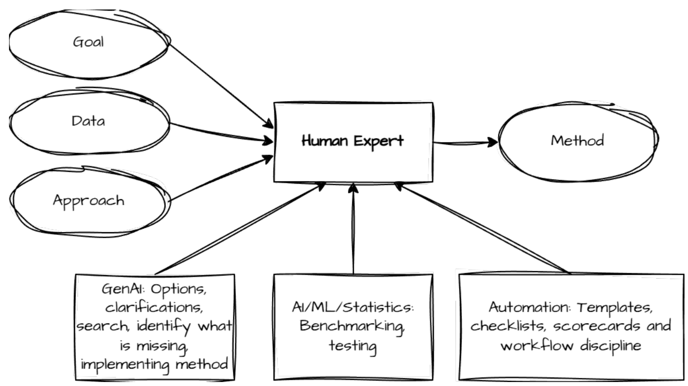

This is another place where local context matters. A method may perform slightly better in theory and still be the wrong choice if it is fragile, opaque, too costly to maintain, or poorly understood by those who need to act on it.

Where GenAI can help

Cognitive Support

GenAI can help as a comparison and critique tool.

Useful roles include:

- explaining candidate methods in plain language,
- comparing methods by assumptions, strengths, limitations, and data requirements,
- mapping business conditions to commonly used methods,
- surfacing obvious mismatches between a method and the situation,
- helping the team understand trade-offs between simplicity, interpretability, and expected performance,
- drafting rationale for method selection,
- critiquing the selected method against known practices and common failure patterns.

This fits our model well. GenAI is useful for broadening awareness and reducing avoidable gaps, but it should not be the one deciding which method the business will trust.

Implementation support

Code-oriented GenAI tools can also help implement and test the selected forecasting method. This may include:

- generating R or Python code for candidate methods
- creating baseline models
- building backtesting routines

- writing evaluation scripts
- producing diagnostic plots and comparison reports
- scaffolding method comparison experiments
- drafting reusable forecasting workflows

This is useful because it can speed up experimentation and reduce mechanical coding effort. However, GenAI does not validate the methodological soundness of what it produces. Generated code may be statistically incorrect, misuse defaults, leak future information, or implement something different from what the team intended. Human review is required to confirm that the implementation is technically correct and faithful to the selected method.

Where other AI methods can help

This is the first step where predictive and analytical methods become materially important.

Useful roles may include:

- benchmarking several candidate forecasting methods,
- evaluating baseline model performance,
- testing sensitivity to seasonality, trend, volatility, or structural change,
- identifying whether segmentation improves forecast quality,
- comparing simple statistical methods with more advanced machine learning approaches,
- estimating whether added method complexity produces meaningful business value.

This is where classical analytics and ML can provide strong evidence. Still, their role is to inform selection, not to remove human judgment from it.

Where deterministic automation is enough

Deterministic automation can help with structure and discipline, for example:

- standard method selection checklists,
- rule-based default method assignment for simple and recurring cases,
- automated baseline comparisons,
- templates for documenting assumptions and fit criteria,
- workflow rules that require comparison against simple baseline methods before approving a more complex one.

This is useful because many organizations jump too quickly to sophistication. Deterministic controls can force a healthier discipline, such as proving that a complex method is actually better than a simpler one.

Risks and failure modes

Common problems in this step include:

- selecting a method that does not fit the data pattern,
- choosing complexity for its own sake,
- selecting a method that cannot be explained or governed well enough,
- using a method that performs well statistically but poorly in practical decision support,
- overfitting the method to historical data,
- allowing GenAI to produce an elegant comparison that hides shallow evaluation,
- failing to compare against simple baselines,
- choosing a method based on team preference rather than business need.

This step is especially vulnerable to tool-driven thinking. Once people start talking methods, they often confuse technical sophistication with business suitability.

Allocation summary

Best fit: Human-led method selection, supported by analytical testing, GenAI for comparison and critique, and deterministic controls to enforce baseline discipline and explicit assumptions.

Step 5: Execute the forecast

What happens in this step

The business applies the selected forecasting method and produces the forecast output. At this step we actually conducting the forecast using the chosen method, and then measuring forecast accuracy by comparing forecast results with actual demand.

Why this step exists

This step exists because all earlier work was preparation. Up to this point, the organization has defined the question, assembled the relevant inputs, chosen the general approach, and selected the specific method. Now it has to run the process and generate a result that can be reviewed and used.

This is also the first step where the forecast becomes something concrete rather than conceptual. That matters, because many weaknesses stay hidden until the method is actually executed.

Type of work

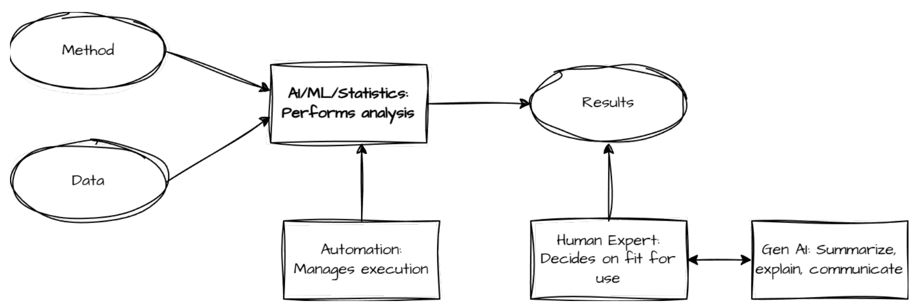

Primarily:
- Prediction work
- Execution work

Secondarily:
- Control work

This is the first step where prediction work is clearly dominant. The process is now doing the computational or judgment-based work needed to produce the forecast.

What requires human judgment

Human judgment is still needed, even though this step is more technical.

It is needed to decide:
- whether the forecast run was executed under the right assumptions,
- whether exceptions, anomalies, or missing inputs need intervention,
- whether the output is plausible enough to proceed,
- whether the result should be accepted as-is, adjusted, or escalated for review,
- whether the observed forecast quality is sufficient for the business purpose,
- whether differences between forecast and actuals reflect model weakness, data problems, or real change in the business.

This is important. Even when execution is heavily automated, the business still needs someone to judge whether the produced result deserves trust.

Where GenAI can help

GenAI can help around the forecast, but usually not as the core engine of execution in a generic demand forecasting process.

Useful roles include:

- explaining forecast outputs in plain language,
- summarizing run results and comparison notes,
- highlighting unusual differences or inconsistencies for review,
- drafting commentary for planners or business stakeholders,
- translating technical output into decision-oriented language,
- helping investigate why a result may look odd by surfacing common causes and diagnostic questions,
- comparing forecast behavior against expected patterns described by the team.

GenAI is especially useful here as an interpreter and critic. It helps make the result easier to understand and review, but it should not be confused with the forecasting method itself unless the forecast design explicitly depends on language-based reasoning.

Where other AI methods can help

This is the center of gravity for predictive and analytical AI.

Useful roles may include:

- running the selected statistical or machine learning forecast,
- generating confidence intervals or uncertainty estimates,
- producing segmented forecasts,
- testing ensemble approaches,
- identifying anomalies during execution,
- comparing actuals to forecast in a structured way,
- supporting automated backtesting and performance evaluation.

If the organization is using quantitative forecasting, this is the step where those analytical methods are doing their main job.

Where deterministic automation is enough

A large part of this step can be handled through deterministic automation, especially in repeatable environments.

Examples include:

- scheduled forecast runs,
- data refresh workflows,
- batch execution pipelines,
- rules for handling missing inputs,
- exception routing,
- generation of standard reports,
- storage of results,
- alerts for forecast deviations or run failures,
- reproducible backtesting routines.

This is an important point. Even when the forecasting method itself is probabilistic, much of the surrounding operational machinery should be deterministic.

Risks and failure modes

Common problems in this step include:

- producing a forecast mechanically without checking whether assumptions still hold,
- accepting output because it looks polished rather than because it is meaningful,
- silent failures in data refresh or execution workflow,

- treating forecast accuracy as the only success criterion while ignoring business usefulness,
- letting GenAI produce persuasive explanations for weak or unstable outputs,
- over-trusting automation and missing obvious signs that the forecast is broken,
- generating a technically correct forecast that arrives too late, at the wrong level, or in the wrong form to support action.

This step is vulnerable to false operational confidence. Once the process runs and generates numbers, people often assume the hard part is done. In reality, this is where badly designed logic becomes visible, and where weak governance can turn output into action too quickly.

Allocation summary

Best fit: Predictive AI or statistical methods for the forecast itself, deterministic automation for repeatable execution and controls, with human review and GenAI support for interpretation and communication.

Step 6: Monitor and adjust the forecast

What happens in this step

The business monitors forecast performance over time and adjusts the method, assumptions, or parameters when needed. This step is a continuous review of forecast accuracy and suitability, with the explicit warning that a method that once worked well may become inappropriate as conditions change.

Why this step exists

This step exists because forecasting is not a one-time act. Demand changes, business conditions change, data quality changes, and the causes of demand may shift. A forecast that was reasonable last quarter may become weak or misleading later. Monitoring exists to detect that drift before bad output quietly turns into bad decisions.

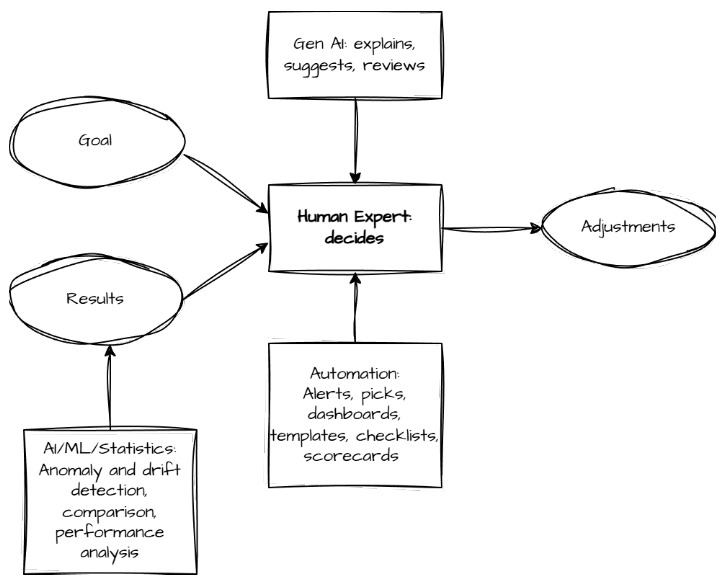

Type of work

Primarily:

- Control work
- Decision work

Secondarily:

- Prediction work
- Knowledge work

This is mainly about checking whether the forecasting process remains fit for purpose and deciding what, if anything, should change.

What requires human judgment

Human judgment is central here.

It is needed to decide:

- whether forecast errors are acceptable in business terms,
- whether a drop in performance is noise, seasonality, a one-off event, or a real structural shift,
- whether the method should be tuned, replaced, or left alone,

- whether the problem is in the model, the data, the business assumptions, or the operating context,
- whether forecast accuracy is the right signal, or whether business usefulness has changed even when statistical metrics still look acceptable,
- whether action should be taken now or whether more evidence is needed.

This is exactly the kind of step where local context matters. A model does not know when the world has changed in a way that matters to this business. It only reflects what it can detect from inputs and metrics.

Where GenAI can help

GenAI can be very helpful here as a reviewer, explainer, and gap-finder.

Useful roles include:
- summarizing forecast performance over time,
- explaining metric changes in plain language,
- comparing recent behavior against known failure patterns,
- surfacing possible causes for deterioration,
- helping structure post-mortem or review discussions,
- drafting recommendations for what to investigate next,
- translating technical monitoring signals into business-facing explanations,
- checking whether the review is missing obvious questions or alternative explanations.

This is a strong fit for GenAI because much of the work is interpretive and comparative. It helps the team reason more broadly and communicate more clearly, but it should not decide by itself whether the forecasting process remains acceptable.

Where other AI methods can help

Other AI and analytical methods can play a significant role here.

Useful roles may include:
- tracking forecast error trends,
- detecting drift or structural breaks,
- identifying anomalies in model performance,
- comparing alternative models over time,
- triggering re-estimation or retraining workflows,
- analyzing whether different segments are degrading differently,
- quantifying uncertainty changes,
- supporting champion-challenger comparisons between current and alternative methods.

This is where predictive and analytical tools help answer, “Is the forecasting system still working well enough?” They can provide evidence, but they do not determine business acceptability on their own.

Where deterministic automation is enough

A large part of monitoring is a very good fit for deterministic automation.

Examples include:
- scheduled accuracy reports,
- threshold-based alerts,
- control charts,
- exception dashboards,
- automated comparisons of forecast versus actuals,

- workflow triggers for review when performance drops below predefined limits,
- logging, audit trails, and change records,
- recurring distribution of monitoring packs to stakeholders.

This is a strong case for deterministic automation. Monitoring should be consistent, repeatable, and visible. AI may help interpret the signals, but the basic control machinery should usually be rule-based.

Risks and failure modes

Common problems in this step include:

- continuing to trust a method long after conditions have changed,
- reacting to random noise as if it were structural change,
- focusing on statistical accuracy while missing declining business usefulness,
- failing to notice that data definitions or source quality have shifted,
- relying on dashboards without real interpretation,
- letting GenAI produce plausible explanations that sound insightful but are not grounded enough,
- changing methods too often and creating instability,
- not changing methods when the evidence clearly calls for it.

This step is vulnerable to two opposite mistakes: complacency and overreaction. Good monitoring is not just seeing the numbers. It is deciding when the numbers mean something important.

Allocation summary

Best fit: deterministic automation for regular monitoring and alerting, analytical methods for drift and performance detection, with human judgment and GenAI support for interpretation, review, and decision framing.

Summary

Main patterns across the demand forecasting process

Across the full process, the main pattern is consistent:

- Human experts define what matters for this business, choose what is good enough, and decide what conclusions and actions make sense.
- GenAI helps reduce narrow thinking and routine effort. It supports framing, comparison, critique, explanation, and, in some steps, technical implementation.
- Other AI, ML, and statistical methods are most useful where the core task is actual prediction, pattern detection, or analytical comparison.
- Deterministic automation is essential wherever the work is repeatable, rule-based, and should be consistent.

The main point is not to maximize AI usage. The point is to assign each kind of work to the form of support that fits it best.

Compact step-by-step view

Step 1. Define what to forecast

Main role of the step:

Define the business question, forecasting target, and planning horizon.

Human expert

Dominant. Decides what matters, what level of forecast is needed, and what “good enough” means.

GenAI

Helps clarify the question, surface options, compare typical approaches, and identify missing considerations.

Other AI/ML/Statistics

Limited role at this point. May help test whether a proposed target is forecastable.

Deterministic automation

Useful for intake templates, required fields, and workflow discipline.

Step 2. Collect relevant input data

Main role of the step

Gather the information needed to support the forecast.

Human expert

Dominant for relevance and sufficiency. Decides which inputs matter and whether the evidence is good enough.

GenAI

Helps identify common input categories, summarize expert input, find gaps, and translate between business and analytics. Also helps with implementation support such as SQL, ETL, preprocessing, and data validation scripts.

Other AI/ML/Statistics

Useful for exploratory analysis, anomaly detection, segmentation, and early signal testing.

Deterministic automation

Very strong role for extraction, transformation, validation, checks, dashboards, and repeatable pipelines.

Step 3. Choose the forecasting approach

Main role of the step

Decide whether the problem is best handled qualitatively, quantitatively, or in hybrid form.

Human expert

Dominant. Chooses the broad logic of attack based on context, horizon, and data reality.

GenAI

Helps compare approaches, explain trade-offs, and critique weak reasoning.

Other AI/ML/Statistics

Helps test whether a quantitative path is viable.

Deterministic automation

Useful for checklists, scorecards, and workflow gates that make rationale explicit.

Step 4. Select the forecasting method

Main role of the step

Choose the specific method within the selected approach.

Human expert

Dominant. Decides what method is appropriate, governable, and worth the complexity.

GenAI

Helps compare methods, explain assumptions, critique mismatches, and document rationale. Also helps with implementation support such as R or Python code, backtesting, evaluation scripts, and experiment scaffolding.

Other AI/ML/Statistics

Strong role in benchmarking candidate methods, testing baselines, and evaluating performance.

Deterministic automation

Useful for baseline comparisons, selection templates, and controls that prevent unjustified complexity.

Step 5. Execute the forecast

Main role of the step

Run the selected method and produce the forecast output.

Human expert

Reviews plausibility, exceptions, and whether the output is fit for use.

GenAI

Helps interpret results, summarize outputs, explain anomalies, and communicate findings.

Other AI/ML/Statistics

Dominant if the forecast is quantitative. This is where the selected method actually does its main job.

Deterministic automation

Strong role for scheduled runs, pipelines, alerts, storage, reporting, and repeatable controls.

Step 6. Monitor and adjust the forecast

Main role of the step

Review forecast performance over time and adjust when needed.

Human expert

Dominant for judgment. Decides whether errors matter, whether conditions changed, and whether intervention is needed.

GenAI

Helps explain performance changes, structure reviews, surface possible causes, and support diagnostic discussions.

Other AI/ML/Statistics

Strong role in drift detection, anomaly detection, champion-challenger comparisons, and performance analysis.

Deterministic automation

Very strong role for alerts, monitoring packs, dashboards, thresholds, and recurring control routines.

Overall role pattern by capability

- Human expert
    - Dominant in Steps 1, 3, 4, and 6
    - Still essential in Steps 2 and 5
    - Owns business meaning, trade-offs, sufficiency, interpretation, and action
- GenAI
- Useful across all steps as cognitive support
    - Most useful for implementation support in Steps 2 and 4
    - Helps reduce routine effort and omission risk
    - Should support decisions, not own them
- Other AI/ML/statistical methods
    - Limited in Step 1
    - Supportive in Steps 2 and 3
    - Strong in Step 4
    - Dominant in Step 5
    - Strong again in Step 6 for performance analysis and drift detection
- Deterministic automation
- Helpful in Step 1
    - *Strong in Step 2*
    - *Helpful in Steps 3 and 4*
    - *Strong in Step 5*
    - *Very strong in Step 6*
    - *Best fit wherever repeatability, consistency, and transparency matter more than flexible reasoning*

	Human Expert	Generative AI	Traditional AI/ML, Statistic, Numeric Analysis	Automation
1. Define what to forecast	**Primary**	Support	*Limitied*	Support
2. Collect relevant input data	**Primary**	**Strong Support**	Support	**Strong Support**
3. Choose Approach	**Primary**	Support	Support	Support
4. Choose Method	**Primary**	**Strong Support**	**Strong Support**	Support
5. Execute Forecast	**Strong Support**	Support	**Primary**	**Strong Support**
6. Monitor and adjust	**Primary**	Support	**Strong Support**	**Strong Support**

Short takeaway

The process follows a clear pattern:

- early steps are mainly about framing and choice
- middle steps are about preparation and method realization
- later steps are about execution, control, and adaptation

That leads to a simple allocation rule:

- keep business-specific judgment with humans
- use GenAI to broaden thinking and reduce implementation friction
- use predictive AI, ML, and statistics for actual forecasting and analytical evidence
- use deterministic automation wherever repeatable control and consistency are enough

When GenAI should not be used

GenAI should not be used as the primary tool when the work depends on business-specific judgment, requires deterministic correctness, or carries a high cost of subtle error.

Its ability to generate plausible language, code, and analysis is useful, but that same strength can also hide weak reasoning, incorrect assumptions, or silent defects. For that reason, there are parts of the demand forecasting process where GenAI should be avoided, tightly bounded, or used only as secondary support.

1. Do not use GenAI to decide what matters for the business

GenAI should not define the forecasting target, the decision purpose, the acceptable level of error, or the action that should follow from the forecast.

These decisions depend on local priorities, constraints, trade-offs, and consequences that are specific to the business. GenAI can help clarify options, but it should not own the choice.

Examples

- deciding whether to forecast by SKU, product family, region, or channel
- deciding whether the forecast is for inventory planning, staffing, budgeting, or strategic investment
- deciding what forecast accuracy is good enough for the business need
- deciding what action should be taken based on the result

2. Do not use GenAI where deterministic correctness is required

GenAI should not be the primary mechanism for tasks that must be exact, repeatable, and auditable.

If the task is rule-based and can be handled by deterministic automation, there should be a strong reason to prefer GenAI.

Examples

- fixed business rules
- standard calculations
- schema validation
- threshold checks
- production ETL logic
- recurring control reports
- workflow gating
- audit trails

In these cases, deterministic automation is usually more reliable, more transparent, and easier to govern.

3. Do not use GenAI as a substitute for actual forecasting methods

GenAI should not replace statistical, analytical, or machine learning methods when the core task is quantitative prediction.

It may help explain, compare, or implement those methods, but it is not the right default tool for numerical forecasting just because it is easy to prompt.

Examples

- estimating future demand from historical time-series data
- comparing baseline forecasting performance
- measuring forecast error over time
- detecting drift or structural change
- evaluating whether model complexity improves business value

These tasks belong primarily to statistics, forecasting methods, and other analytical techniques.

4. Do not use GenAI without strong review when the output looks polished but is hard to verify

GenAI is especially risky when it produces something that sounds convincing, reads well, or runs without error, but is difficult for the team to validate.

This includes both language output and generated code.

Examples

- a polished problem statement that hides ambiguity
- a method comparison that sounds informed but skips critical assumptions
- generated R or Python code that runs but implements the wrong method
- backtesting logic that leaks future information
- summaries that smooth over weak source material
- explanations that feel insightful but are not grounded enough

The danger here is not visible failure. The danger is false confidence.

5. Do not use GenAI as the final judge of completeness or readiness

GenAI should not be the final authority on whether the work is finished, whether the method is appropriate, or whether the output is fit for business use.

It can critique and help review, but it should not close the loop by itself.

Examples

- deciding whether the data is sufficient
- deciding whether the selected method is acceptable
- deciding whether forecast quality is good enough
- deciding whether a decline in performance is important
- deciding whether business action should proceed

These are judgment-heavy decisions and should remain human-led.

6. Do not use GenAI where the cost of a subtle mistake is high and validation is weak

GenAI becomes especially dangerous when:

- the error may not be obvious,
- the result may still look reasonable,
- and the business impact of being wrong is material.

In those situations, GenAI may still help as a secondary tool, but it should not be trusted as the main mechanism.

Examples

- demand forecasts tied directly to inventory commitments
- forecasts used for staffing or capacity decisions with high cost of error
- executive reporting where weak reasoning can shape major decisions
- code generation for production forecasting pipelines without expert review

7. Do not use GenAI just because the task feels intellectual

Some tasks look like "AI work" only because they involve text, explanation, or complexity. That does not mean GenAI is the right answer.

A useful test is simple:

- if the task requires business judgment, keep it human-led
- if it requires exact repeatability, prefer deterministic automation
- if it requires quantitative prediction, use analytical methods

- use GenAI when it improves thinking support or reduces implementation friction without taking ownership away from the right actor

Practical rule

GenAI should not be used as:

- the owner of the business question
- the owner of the forecasting decision
- the substitute for deterministic controls
- the substitute for forecasting methods
- the final judge of correctness, completeness, or business readiness

It is most useful when it supports the work, not when it silently takes control of it.

Short takeaway

Do not use GenAI where the work must be:

- specific to this business,
- exactly correct,
- quantitatively grounded,
- or explicitly accountable.

Use it to broaden thinking, reduce routine effort, and accelerate implementation where the intent, method, and review are already clear.

Contents