

Practice Guide No. 004

The Predictive Analytics Readiness Review

A CPA-oriented worksheet for deciding whether a predictive use case is ready for deeper modeling, validation, or implementation review

Valentina DuPont, CPA | Knowledge Hub | Release 004 | Published August 10, 2026 | Version 1.0

Purpose: Help finance and accounting professionals evaluate whether a predictive analytics idea is sufficiently defined, supported, and governed to move forward. This guide does not certify a model for deployment and does not replace qualified statistical, legal, compliance, data-science, or model-risk expertise.

How to Use This Guide

Use the guide before a predictive initiative is funded, piloted, materially relied upon, or expanded. For each section, mark the status that best reflects the evidence available today, note unresolved questions, and identify where specialist review is required.

Clear - the question/evidence is sufficiently defined for the current stage.

Needs work - the issue is understood but not yet resolved.

Specialist input - deeper technical, legal, compliance, data, or model-risk expertise is required.

Not applicable - document why the item does not apply.

1. Business Question & Decision

☐ Clear ☐ Needs work ☐ Specialist input ☐ N/A

☐ What exact outcome are we trying to predict?

☐ What decision will change because of the prediction?

☐ Who owns that business decision?

☐ What prediction horizon and decision date matter?

☐ What would success look like in business terms?

☐ Is prediction actually necessary, or could a simpler rule/process solve the problem?

Notes / evidence / specialist follow-up:

2. Historical Data & Evidence

☐ Clear ☐ Needs work ☐ Specialist input ☐ N/A

☐ What population and time period will be used?

☐ Are definitions consistent across periods, systems, and business units?

☐ Was every proposed input actually available at the prediction date?

☐ Are missing values, duplicates, labels, joins, mappings, and category definitions understood?

☐ Do any features depend on judgment, third-party data, generative AI, or alternative data?

☐ Can the final output be traced back to source, definition, timing, and transformation?

Notes / evidence / specialist follow-up:

3. Capability & Roles

☐ Clear ☐ Needs work ☐ Specialist input ☐ N/A

☐ Do we have the statistical/modeling expertise needed for this use case?

☐ Do we need a data engineer, data scientist, econometrician/statistician, domain expert, model-risk specialist, or legal/compliance reviewer?

☐ Who prepares the model and who independently challenges it?

☐ Who can change data definitions, model parameters, or decision rules?

☐ Are role boundaries and escalation paths documented?

Notes / evidence / specialist follow-up:

4. Validation & Performance Evidence

- ☐ Clear ☐ Needs work ☐ Specialist input ☐ N/A
- ☐ What baseline must the model beat?
- ☐ How will training, validation, and test data be separated?
- ☐ For time-dependent data, does testing preserve chronology and avoid future information?
- ☐ What evidence will show that the model generalizes beyond the fitting sample?
- ☐ Which performance metrics will be used, and what does each metric actually prove?
- ☐ If exact probabilities drive decisions, have we considered calibration as well as ranking/discrimination?
- ☐ What result would cause us to reject, redesign, or stop the model?

Notes / evidence / specialist follow-up:

5. Interpretation, Prediction & Causation

- ☐ Clear ☐ Needs work ☐ Specialist input ☐ N/A
- ☐ Are we using the model to predict what is likely to happen, or to justify changing something?
- ☐ Could a predictive variable be an indicator rather than a cause?
- ☐ What additional evidence is needed before acting as though a relationship is causal?
- ☐ Are limitations and uncertainty visible to the decision-maker?

Notes / evidence / specialist follow-up:

6. From Prediction to Action

☐ Clear ☐ Needs work ☐ Specialist input ☐ N/A

☐ What exactly does the model produce: a number, probability, score, ranking, or category?

☐ How does that output become a recommendation or action?

☐ What thresholds, optimization rules, or business logic sit between prediction and decision?

☐ What actions are allowable, prohibited, or subject to approval?

☐ Who has final authority, and how are overrides handled and documented?

Notes / evidence / specialist follow-up:

7. Economics, Adoption & Operational Fit

☐ Clear ☐ Needs work ☐ Specialist input ☐ N/A

☐ Does expected value exceed the total cost of data, technology, implementation, validation, training, monitoring, and maintenance?

☐ Can the recommendation be integrated into the actual workflow?

☐ Will users understand when and how to rely on it?

☐ How will exceptions and disagreements be handled?

☐ What evidence will show that the model improves the current process rather than merely performing well statistically?

Notes / evidence / specialist follow-up:

8. Governance, Fairness & Lifecycle

☐ Clear ☐ Needs work ☐ Specialist input ☐ N/A

☐ Are any variables legally restricted, sensitive, or potential proxies for protected characteristics?

☐ Are similarly situated customers or transactions treated consistently?

☐ What controls constrain optimization or automated action?

☐ Who monitors model performance, data drift, calibration, exceptions, and outcomes?

☐ What changes trigger recalibration, retraining, escalation, suspension, or retirement?

☐ Who remains accountable for the consequences of the decision system?

Notes / evidence / specialist follow-up:

Readiness Decision

Select the best current disposition. This is a planning decision, not a deployment certification.

☐ Not ready - The business question, data, evidence, ownership, or control environment is not sufficiently defined to proceed responsibly.

☐ Discovery ready - The use case is worth deeper investigation, but material data, technical, economic, or governance questions remain.

☐ Pilot ready - with appropriate specialists - The question and initial evidence are sufficiently defined for a controlled pilot with qualified technical and governance support.

☐ Candidate for formal implementation review - A controlled pilot has produced enough evidence to justify a separate, formal implementation and model-governance review. This guide does not itself authorize deployment.

Decision rationale

One-Page Readiness Checklist

- ☐ Business outcome and decision are defined precisely.
- ☐ Historical data is complete enough, consistently defined, and available at the correct decision time.
- ☐ Key mappings, categories, third-party inputs, and AI-generated assumptions are traceable and reviewed.
- ☐ Required technical and domain specialists are identified.
- ☐ Training/test design reflects real use and avoids leakage or future information.
- ☐ Performance metrics are understood, including what they do not prove.
- ☐ Prediction is not being confused with causation.
- ☐ Decision rules, allowable actions, approvals, and override paths are documented.
- ☐ Economics and operational adoption have been assessed, not just statistical performance.
- ☐ Fairness, legal, ethical, and policy constraints are addressed where relevant.
- ☐ Monitoring, change control, escalation, retraining, and retirement criteria are defined.
- ☐ A named human owner remains accountable for the business outcome.

A strong predictive model is evidence. A dependable business system requires the surrounding data, economics, decision rules, controls, and accountability to work as well.

Learning Note / Disclaimer

This resource is an educational decision-support worksheet developed for the Valentina DuPont, CPA Knowledge Hub. It is not investment, legal, regulatory, statistical, model-validation, accounting, or other professional advice. Predictive systems may require qualified specialists and organization-specific review before use in consequential decisions.