

WHITE PAPER

PRAGMATIC PROBABILISTIC ESTIMATING (PPE)

A Structured Alternative to Single-Point Guessing

A practical framework that combines expert judgement, historical evidence and probabilistic thinking to produce better estimates and better conversations about uncertainty.



STRUCTURED
EXPERT INPUT



PRACTICAL
PROBABILISTIC
INSIGHT



BETTER
CONVERSATIONS.
BETTER DECISIONS.



A PRAGMATIC PROBABILISTIC ESTIMATING FRAMEWORK: A STRUCTURED ALTERNATIVE TO SINGLE-POINT GUESSING

Sean Whitaker
sean@crystal.consulting
Crystal Consulting
May 2026

INTRODUCTION

Before anybody immediately criticises this framework because it is not full Reference Class Forecasting, Monte Carlo simulation, Bayesian modelling, or advanced quantitative risk analysis, they should first pause long enough to understand what this paper is actually trying to solve. I get that you really favour lots of science and maths, but not everyone does or has access to the tools or knowledge to do that. This framework is not attempting to compete with mature probabilistic analysis performed by experienced statisticians using large, clean datasets and sophisticated simulation models. Organisations that possess that level of capability should absolutely continue using it where appropriate (and they can afford it) because full probabilistic modelling remains one of the strongest approaches available for understanding uncertainty in complex projects.

The problem is that many organisations are nowhere near that level of maturity. In reality, a large proportion of projects across infrastructure, technology, engineering, operational improvement, construction, and organisational change are still being estimated using simplistic deterministic methods (or guesstimates disguised as objectivity). Somebody is asked for a number. They provide a number. That number is then discussed, negotiated, reduced, rounded, approved, and eventually treated as though it represented certainty, when in fact it usually never did.

Many organisations know their estimating practices are weak, but they also lack the data, capability, budget, governance maturity, or time required to perform sophisticated quantitative analysis. They exist in an uncomfortable middle ground where simplistic estimating is no longer good enough, but advanced modelling remains impractical. This model proposes a practical bridge across that gap.

The framework described here combines several established concepts including historical evidence, Delphi estimating, structured expert judgement, facilitated discussion, three-point estimating, and lightweight probabilistic interpretation. The purpose is not to eliminate uncertainty because uncertainty cannot be eliminated. The purpose is to expose uncertainty more honestly, structure estimating conversations more effectively, reduce obvious forms of bias, and provide decision makers with something substantially better than a single-point guess pretending to be a fact.

Most importantly, this framework is intended to be practical. An organisation should be able to implement it immediately using spreadsheets, forms, workshops, facilitation, and relatively basic analytical capability. No specialist software is required, no statistician is required., and no advanced mathematical background is required to begin improving estimating maturity.

This framework is most effective in environments where uncertainty is high enough that deterministic single-point estimates are unreliable, but where the organisation does not yet possess the capability, historical datasets, time, budget, or organisational maturity required for full probabilistic modelling. It is particularly useful for projects involving significant expert judgement, ambiguity, evolving scope, organisational change, software development, integration complexity, innovation, or uncertain delivery conditions where precise quantities, production rates, or productivity factors are not yet fully understood. Conversely, this framework is generally not the best choice for highly repetitive or well-understood work where strong parametric estimating can be confidently applied using known quantities, established productivity rates, mature historical datasets, and statistically reliable production information. For example, if an organisation already knows with reasonable confidence the cost per metre of pipe installation, the productivity rate for laying asphalt, or the average effort required per standard software transaction, then parametric estimating will often provide stronger and more defensible forecasting. The Pragmatic Probabilistic Estimating (PPE) framework is intended primarily for the large middle ground where projects are too uncertain for simplistic deterministic estimation, but not mature enough to justify full quantitative probabilistic analysis.

At an absolute minimum, this framework creates better conversations, and in many organisations, better conversations about uncertainty, assumptions, confidence, contingency, and delivery realism would represent a major improvement by themselves.

THE REAL PROBLEM WITH PROJECT ESTIMATING

One of the biggest misconceptions in project management is that estimating is primarily a mathematical problem. In reality, estimating is usually a behavioural and organisational problem long before it becomes a statistical one. Most organisations do not fail because they lack advanced probability modelling capability. They fail because the conversations surrounding estimates are weak, politically distorted, incomplete, or fundamentally dishonest about uncertainty.

Many estimates are not actually estimates at all. They are targets disguised as forecasts. or they are numbers selected because they sound commercially acceptable, politically palatable, or strategically convenient. The person providing the estimate may already know the number is optimistic (but usually doesn't) but understands that presenting a realistic range could create organisational resistance, approval delays, or accusations of negativity – meaning the project doesn't get the go ahead. As a result, uncertainty becomes suppressed very early in the lifecycle.

The project then proceeds under an illusion of precision. Stakeholders begin making decisions as though the estimate represents something factual and reliable, when in reality it may contain enormous hidden variability. Eventually reality catches up and some or all of the following occurs:

- scope expands,
- risks emerge,
- dependencies fail,
- procurement takes longer than expected,

- integration complexity appears,
- stakeholders change direction,
- technical uncertainty materialises,
- the estimate begins to unravel

At that point, organisations often behave as though the project team failed to deliver accurately. In truth, the estimate may never have honestly represented the level of uncertainty that existed from the beginning. This is one of the central ideas underpinning this framework. Projects are uncertain by nature. Pretending otherwise does not reduce uncertainty. It merely delays its appearance until later in the project lifecycle when it becomes more expensive, more disruptive, and more politically visible.

WHY MOST ORGANISATIONS NEED A MIDDLE GROUND

At one end of the estimating maturity spectrum are low-maturity approaches dominated by intuition, single-point estimates, rough-order-of-magnitude assumptions (or estimates), and executive pressure. At the other end are highly sophisticated approaches such as Monte Carlo simulation, Reference Class Forecasting, Bayesian analysis, stochastic modelling, and advanced quantitative risk analysis. The problem is that many organisations are somewhere in the middle.

They are mature enough to recognise that simplistic estimating is unreliable, but not mature enough to support full probabilistic modelling. They may not possess:

- reliable historical datasets,
- experienced quantitative analysts,
- mature governance capability,
- advanced modelling tools,
- or sufficient project scale to justify sophisticated analysis.

And even if they did, many organisations would still struggle culturally to interpret and use those outputs correctly. This framework deliberately acknowledges that reality. It is not designed for organisations already operating at the highest levels of quantitative forecasting maturity. It is designed for organisations attempting to move beyond simplistic deterministic estimation without needing to become statistical modelling experts overnight. In many respects, this framework is less about mathematics and more about improving organisational estimating behaviour and conversations around uncertainty.

WHAT THIS FRAMEWORK ACTUALLY DOES

At its core, this framework attempts to achieve five things simultaneously.

- First, it attempts to anchor estimating conversations in reality wherever possible by introducing historical evidence into the process.
- Second, it attempts to reduce obvious forms of bias by using anonymous structured estimation rather than immediately relying on group workshops dominated by hierarchy and personality.

- Third, it deliberately forces participants to think in ranges rather than deterministic numbers by requiring optimistic, most likely (or realistic), and pessimistic estimates.
- Fourth, it introduces lightweight probabilistic thinking without requiring advanced simulation capability.
- Finally, it attempts to improve governance conversations by reframing estimates as uncertainty-informed forecasts rather than commitments pretending to be facts.

This distinction matters enormously as the framework is not a prediction engine, it is instead a structured decision-support framework.

HOW TO IMPLEMENT THE FRAMEWORK

The framework itself is relatively straightforward to implement operationally, although the quality of outputs will depend heavily on discipline, facilitation quality, and organisational honesty.

The process begins before any estimating occurs as the first requirement is to clearly define exactly what is being estimated. This sounds trivial, but poor estimating often begins with scope ambiguity. Participants cannot produce meaningful estimates if they are unknowingly estimating different things.

If estimating a project, the organisation should clearly define what is included and excluded from the estimate. For example, does the estimate include procurement activities? User training? Integration testing? Hypercare support? Vendor delays? Data migration? Regulatory approval? Infrastructure upgrades? Operational transition?

These details matter because different participants will otherwise unconsciously make different assumptions. Once scope clarity has been established, the organisation should gather whatever historical information exists. Importantly, the framework does not require perfect historical datasets. Many organisations will not possess them. Even partial historical information can still provide valuable calibration.

This information may include comparison of initial estimates and actual durations from previous projects, initial estimates and actual costs, effort hours, procurement lead times, defect rates, change volumes, productivity metrics, or lessons learned. The goal here is not to treat historical information as deterministic truth. The goal is to force realism into the conversation.

For example, if previous projects of comparable complexity consistently took fourteen to sixteen months, and the current estimate being proposed is eight months, the organisation should immediately ask a simple question:

“What is different this time that justifies such a substantial reduction?”

That single question often reveals hidden optimism, incomplete scope understanding, political pressure, or unrealistic assumptions.

Once the available historical information has been assembled, the organisation begins the first Delphi estimation round. Participants should not simply be asked to provide “the

estimate.” That almost guarantees a simplistic deterministic response. Instead, participants should be asked to provide three separate estimates:

- an optimistic estimate,
- a most likely (or realistic) estimate,
- and a pessimistic estimate.

The wording used here is extremely important. Participants should understand that optimistic does not mean fantasy. It means a genuinely achievable outcome if conditions are unusually favourable. Likewise, pessimistic does not mean catastrophic failure. It means a realistic outcome if significant but plausible issues arise. Participants should also be asked to explain:

- the assumptions underpinning their estimate,
- major risks,
- dependencies,
- constraints,
- and areas of uncertainty.

Crucially, this process should be anonymous. This is one of the most important design features in the framework because it helps reduce authority bias, hierarchy effects, groupthink, political influence, anchoring behaviour, and dominant personalities.

In many traditional workshops, the most senior person speaks first and unintentionally anchors the entire discussion. Junior participants often modify their estimates toward what appears politically acceptable rather than what they genuinely believe. The anonymity in the Delphi process disrupts that behaviour.

Participants complete their estimates independently using forms, spreadsheets, surveys, or simple digital tools. The facilitator then consolidates the responses without identifying contributors. At this stage, the organisation should resist the temptation to immediately force consensus. Wide estimate variation is often extremely valuable information. If one participant estimates six months and another estimates eighteen months, that difference likely indicates fundamentally different assumptions, risk perceptions, or understanding of scope. The disagreement itself becomes useful intelligence.

The facilitator then redistributes the anonymous responses to all participants. This redistribution should include estimate ranges, assumptions, rationale, recurring themes, and major risks identified by participants. Participants are then asked whether they wish to revise their estimates after reviewing the thinking of others.

This stage is one of the most intellectually valuable aspects of the framework because it allows participants to reconsider assumptions without embarrassment or political exposure. People often improve their judgement when exposed to additional information. The anonymity simply allows them to do so honestly. In many cases, the estimate range naturally narrows during this iterative process. If substantial disagreement still exists after the redistribution round, the organisation should then conduct a facilitated workshop. Importantly, the purpose of the workshop is not to force agreement, its purpose is to understand disagreement.

The facilitator should deliberately explore why estimates differ. What assumptions are different? Which risks are being weighted differently? Which constraints are interpreted differently? What uncertainties are driving the divergence?

The quality of facilitation matters enormously here. Poor facilitation rapidly reintroduces hierarchy effects and political distortion. Strong facilitation ensures all perspectives are heard while keeping the discussion evidence-focused rather than personality-driven.

Following the workshop, participants should then complete another anonymous estimation round. This is critically important because people frequently change their thinking after hearing additional information but may hesitate to publicly revise their earlier positions. The second anonymous round allows participants to recalibrate privately and as a result the resulting estimates are usually substantially stronger than the initial submissions.

CONVERTING THE ESTIMATES INTO PRACTICAL PROBABILISTIC OUTPUTS

Once the organisation has final optimistic, most likely, and pessimistic estimates, the framework then introduces lightweight probabilistic interpretation.

The traditional PERT weighted average formula is used as a practical approximation:

$$E = \frac{(O + (4xM) + P)}{6}$$

Where:

- O represents the optimistic estimate,
- M represents the most likely estimate,
- and P represents the pessimistic estimate.

The formula deliberately weights the most likely estimate more heavily than the optimistic and pessimistic estimates. For example, if:

- the optimistic estimate is 8 months,
- the most likely estimate is 12 months,
- and the pessimistic estimate is 20 months,

then the weighted expected estimate becomes:

$$E = \frac{8 + (4x12) + 20}{6}$$

$$E = \frac{8 + 48 + 20}{6}$$

$$E = \frac{76}{6}$$

$$E = 12.67$$

This produces an expected estimate of approximately 12.7 months.

Importantly, the framework openly acknowledges that this is not statistically rigorous forecasting. The formula itself relies on assumptions about distribution behaviour that may not accurately reflect real project environments. For example, many projects exhibit skewed distributions, nonlinear escalation effects, and fat tails, particularly in software, infrastructure, organisational change, and highly integrated delivery environments. Nevertheless, this approach still introduces a far more mature conversation than deterministic single-point estimation.

The ultimate goal of any maturing organisation would be to gather data about its performance to develop its own dataset and then adjust the weightings of each of the 3 elements to reflect its actual performance. But in lieu of that data this is as good a place as any to begin.

The framework then introduces a lightweight approximation of standard deviation using:

$$SD = \frac{P - 0}{6}$$

Using the earlier example:

$$SD = \frac{20 - 8}{6}$$

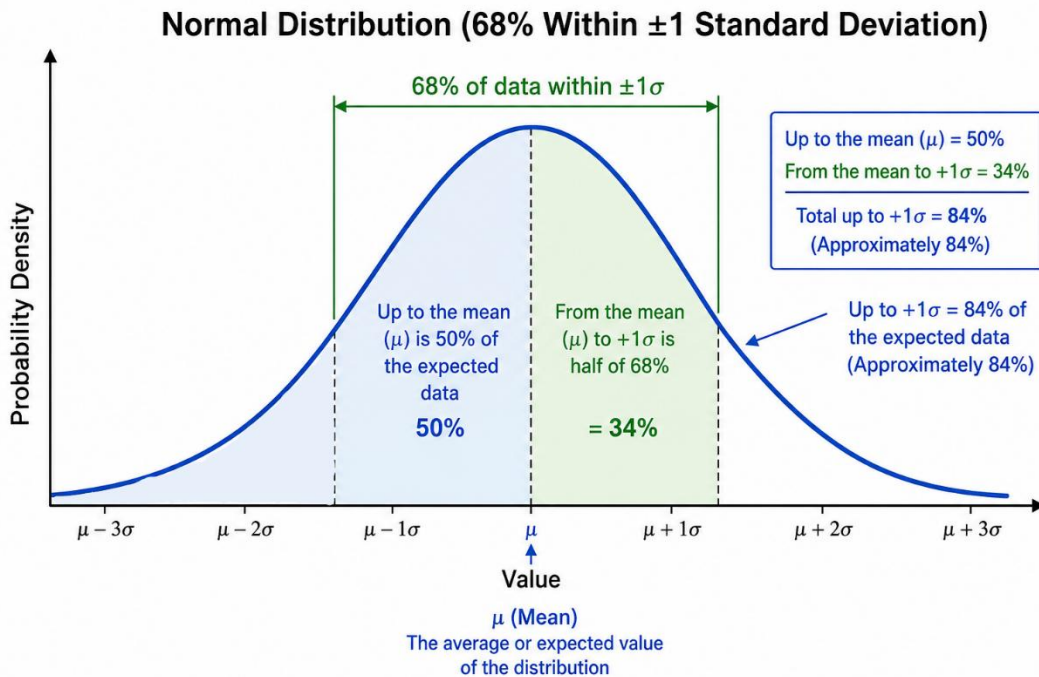
$$SD = \frac{12}{6}$$

$$SD = 2$$

This produces an approximate standard deviation of 2 months. Again, this is not presented as rigorous statistical modelling. It is a practical approximation designed to support organisational discussion around variability and uncertainty.

THE APPROXIMATE P84 DISCUSSION POINT

One of the most practically useful outputs from this framework is the approximate P84 value. Under a roughly normal distribution assumption, one standard deviation above the mean corresponds approximately to an 84% confidence level. This can be represented as:



Using the earlier example:

- the mean estimate is 12.7 months,
- and the approximate standard deviation is 2 months.

The approximate P84 therefore becomes the mean plus 1 standard deviation:

P84 (approx.) $12.7 + 2 = 14.7$ months

This means, approximately speaking, that there is an 84% chance of delivering the project at, or under, 14.7 months, and conversely there is still a 16% chance of the organisation taking longer than 14.7 months. This creates an extremely useful governance discussion point, and the organisation can now begin discussing whether:

- the contingency is sufficient,
- stakeholders are comfortable with the confidence level,
- the risk exposure is acceptable,
- or additional mitigation activities are warranted.

However, this distinction must be made clear (again): This approximate P84 is NOT equivalent to a statistically rigorous P80 or P84 generated through full Monte Carlo simulation or mature probabilistic analysis. A true P80 confidence level derived from advanced simulation would require:

- large-scale scenario simulation,
- dependency modelling,
- correlation analysis,
- distribution fitting,
- and substantially more sophisticated mathematics.

This framework does not claim to replace that capability. Instead, the approximate P84 should be viewed as:

- a lightweight probabilistic approximation,

- a practical confidence discussion mechanism,
- and a significant maturity improvement over deterministic guessing.

That distinction matters enormously because the framework should never create false confidence about the precision of its outputs.

THE REAL VALUE OF THE FRAMEWORK

Ironically, the greatest value of this framework may not be the mathematical outputs at all. The greatest value may simply be that it forces organisations to have much better conversations because the framework, and the outputs from it, encourages discussion around:

- assumptions,
- uncertainty,
- confidence,
- historical performance,
- risk exposure,
- variability,
- contingency,
- optimism bias,
- and delivery realism.

Many organisations never discuss these things openly during estimating and this framework deliberately surfaces them. It helps organisations stop pretending estimates are facts and start treating them as informed forecasts generated under uncertainty. That shift alone represents a major maturity improvement.

CONCLUSION

Projects are uncertain by nature, and no estimating framework can eliminate that uncertainty. What organisations can improve, however, is the quality of their thinking, the honesty of their conversations, and the maturity of their decision-making.

This framework provides a practical and achievable way to move beyond simplistic deterministic estimation without requiring advanced statistical capability. It introduces structured expert judgement, historical evidence, anonymous refinement, probabilistic thinking, and facilitated uncertainty discussion into a single practical process.

It is not a replacement for mature quantitative risk analysis where such capability exists and is justified, but for many organisations, it may represent something just as important: A practical, evidence-informed bridge between guesswork and true probabilistic thinking.