


BEYOND BETTER MODELS

Predictive Analytics Won't Save a Programme That Refuses to Listen

On the distance between what our forecasting models can see and what our governance will
act on



Giovanni Leonardi

November 2005

Predictive Analytics Won't Save a Programme That Refuses to Listen

Programme · No. 43

Originally written November 2005 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The binding constraint on predictive analytics is not analytical; it is temperamental.”

1. The Forecast That Everyone Believed and No One Used

Three quarters into a programme I will not name, the earned value figures told a single, unambiguous story. The schedule performance index had sat below 0.85 for two consecutive reporting periods. The cost performance index was drifting down through 0.9. The estimate at completion, extrapolated from the programme's own actuals rather than from anyone's hopes, put delivery some nine months beyond the committed date and close to a fifth above the approved budget. None of this was contested. The figures were on the screen; the method was orthodox; every head in the room could follow the arithmetic from first principles.

The following week, the board pack went out with delivery confidence rated amber-green and a narrative explaining that the fourth quarter would recover the position.

I have come to treat that sequence not as an aberration but as the central fact about predictive analytics in programme delivery. The model was not wrong. In my experience the model is almost never wrong in the way the vendor demonstrations warn it might be — a decimal place out, a mis-keyed baseline. It was wrong about nothing. What failed was everything that happened after the number reached the table.

2. What We Were Promised

The pitch is a good one, and I want to state it at its strongest before taking issue with it. Instrument the programme properly — earned value, resource actuals, a live risk register, a defect log — feed that data into a forecasting engine, extrapolate the trend or simulate the outcome ten thousand times, and the programme's future becomes visible while there is still time to change it. The forecast, being arithmetic, is neutral. It does not care whose reputation is attached to the committed date. It replaces the soft politics of optimism with the hard discipline of quantification, and it does so early, when intervention is still cheap.

There is real merit in this. A programme that forecasts is better run than one that does not, and the analytical machinery of the last decade — earned value management taken past simple variance reporting, Monte Carlo simulation of schedule risk, trend analysis on the cost curve — is a genuine advance on the wetted finger held up to the wind. I have no quarrel with the mathematics. My quarrel is with the belief that the mathematics is the hard part.

The promise of predictive analytics is that a good enough model will see the future. The reality is that most programmes could already see the future perfectly well, and chose not to look.

3. Where the Promise Meets the Organisation

Three things break the promise, and none of them is a modelling problem.

The data is motivated, not neutral. A forecast inherits the honesty of its inputs, and the most consequential input in programme delivery — percent-complete — is very often an estimate supplied by the very people whose progress it purports to measure. When a workstream lead reports a task ninety per cent finished, they are not lying; they are optimistic, and they are reporting into a governance structure that has made plain it prefers good news. Earned value built on that number is not neutral arithmetic. It is optimism, compounded and given a decimal point. A model fed on comfortable data forecasts comfortably, and it is precisely the least comfortable inputs — the workstream everyone avoids discussing — that it is least likely to receive intact.

The forecast is most inconvenient at the exact moment it becomes useful. A prediction that confirms the plan tells you nothing you did not already believe. The only forecast worth having is the unwelcome one — and the unwelcome one arrives into an organisation exquisitely equipped to neutralise it. There is always a reason the trend will recover: the new supplier ramps up next quarter, the back end always catches up towards the end, the first-quarter numbers were distorted by mobilisation. Where those explanations run out, the baseline itself can be moved, and a re-baselined programme is, by construction, back on plan. The machinery that ought to act on the forecast is the same machinery with the strongest motive to explain it away.

The precision is false where it matters most. A single-point estimate at completion, quoted to the nearest thousand, implies a confidence the underlying data cannot support. A Monte Carlo S-curve offering a delivery date at the eightieth percentile is heard around the table as an eighty per cent chance of success, when what it actually represents is the aggregation of the variances that someone thought to model. Simulation is superb at combining many independent, familiar uncertainties. It is blind to the correlated, structural ones that actually sink programmes: the prime supplier that walks, the regulatory ruling that lands the wrong way, the systems integration everyone privately knew was the real risk and no one entered on the register. Models are confident about the kind of trouble that resembles the past.

4. The Objection I Take Seriously

The natural reply is that these are all solvable problems. Better data quality at source. Independent estimators who do not report to the programme they assess. And, most promising of all, the "outside view" — forecasting a programme not from its own bottom-up plan but by comparing it against the actual, finished outcomes of a reference class of similar programmes, on the reasoning that your programme is far less unique than its sponsors believe. I think this last idea is the most important development in the field, precisely because it attacks the optimism at its source: it takes the estimate out of the hands of the optimists altogether.

I hold no brief against any of these. They move the numbers closer to the truth, and closer to the truth is worth having. But notice what none of them can manufacture. None of them supplies the organisation's willingness to believe an unwelcome figure and act on it while acting is still cheap. You can drive the forecast error down to a rounding error and change nothing at all, if the meeting that receives the forecast has already decided what it intends to conclude. The binding constraint on predictive analytics is not analytical; it is temperamental. We have become fluent in producing the number and remained illiterate in accepting it.

5. What Actually Changes the Outcome

If the constraint is organisational, then that is where the effort belongs. The programmes I have watched use forecasting well were not the ones with the most sophisticated models. They were the ones that had done the unglamorous work of making an unwelcome forecast impossible to ignore.

- Separate the estimator from the estimated. The person who forecasts completion should not be the person whose reputation rests on the committed date. This is not an accusation of dishonesty; it is a recognition that no one marks their own homework severely.
- Watch the forecast's own trajectory, not only its latest value. An estimate at completion that moves outward every period is telling you something more reliable than any single reading: that the programme does not yet understand its own remaining work. A steadily drifting figure is itself the leading indicator.
- Pre-commit to the trigger points. Agree, in advance and in writing, that a schedule performance index below a set threshold for two periods compels a defined intervention — a re-planning gate, an independent review — rather than a discretionary conversation about whether matters are really that serious. Discretion is where inconvenient forecasts go to die.
- Read the indicators that move first. Cost and schedule are lagging measures; by the time they turn, the cause is months old. Rework rates, the time taken to close open decisions, unplanned staff turnover, the defect trend — these move before the earned value does, and a practitioner who watches them is forecasting in a way no cost model can.
- Protect the person who raises the alarm early. An organisation that celebrates the heroic late recovery and quietly marks down the accurate early warning is training its people to withhold precisely the data the model needs. The incentives around the forecast matter more than the forecast.

6. The Judgement Was Never Removed

The enduring promise of predictive analytics was that it would take the judgement out of forecasting — that enough data and a good enough model would settle the question of where a programme was heading and spare us the argument. It has not done that, and I have come to think it never will, because it was never really the arithmetic we were arguing about.

What analytics does, at its best, is relocate the judgement. It moves the hard question from what will happen — which the numbers can now answer well enough — to whether we are

willing to believe them — which no model has ever answered for anyone. That second question is a matter of organisational character, not computation. The forecast on the screen three quarters into that programme was, in the end, correct in every particular. The failure was not that we could not see what was coming. It was that seeing it was never the part we found hard.