Robust Decision Making: From Probability to Possibility

We have built our world on models. As engineers, scientists, and decision-makers, we rely on them to simplify complexity, predict outcomes, and guide our choices. The approach is intuitive: define a system mathematically, identify controllable parameters (levers, *L*), and optimise for a desired outcome (measures, *M*). For well-understood systems – such as engineering designs or physical processes – this method works effectively, as both the system's behaviour and its influencing factors are known. However, as the scope of modelling expanded into less predictable domains – economies, climate systems, or social dynamics – its limitations became apparent.

The problem is not modelling itself, but its overapplication. Optimisation, by definition, eliminates redundancy, leaving systems brittle. This fragility is exacerbated in a *VUCA* world (volatile, uncertain, complex, ambiguous) [1], a term originally coined by the US military to describe modern operational environments. Ironically, optimisation-driven decision-making has contributed to this instability, as cost-efficiency often comes at the expense of robustness.

Traditional responses to modelling failures have followed predictable patterns. The most obvious, especially in academic circles, is simply to build better models: identifying where existing frameworks fall short and proposing more sophisticated ones to fill the gap. J. Doyne Farmer [2] from Oxford, for instance, argues that with sufficient investment, we could model the global economy down to the individual level, and thereby prevent future economic crises. His reference point is weather forecasting – undoubtedly a triumph of applying enormous computational effort to an extraordinarily complex system, at an annual cost of around 5 billion USD. It's a compelling analogy. But one has to wonder: weather forecasters are happy to call a two-week horizon a success. Would Farmer settle for the same?

Other methods attempt to incorporate uncertainty directly into the optimisation. Stochastic optimisation acknowledges that certain parameters are unknown – but then fixes that problem by estimating yet another unknown parameter: the standard deviation. Aside from the logical awkwardness of addressing a lack of knowledge with more estimated knowledge, Nassim Taleb [3], points out the deeper flaw: the most consequential events in our world are precisely those that sit in the tail of the bell curve – low probability, massive impact – and are therefore systematically dismissed by this approach. Vulnerability analysis takes a different angle, examining how an optimal result shifts when individual parameters deviate from their estimates. It's a valid idea, but typically only varies one parameter at a time, leaving the vast space of simultaneous deviations unexplored. Scenario analysis goes further, constructing a range of plausible futures to plan against – capturing the simultaneous shift of multiple parameters – but usually results in a handful of hand-crafted scenarios, chosen by analysts based on experience rather than data.

What all these approaches share is a subtle but damaging tendency: they shift the conversation away from decisions and toward assumptions. When a model is trusted to make the call, the debate naturally gravitates to what's in the model and how it's specified. In a complex world where perfect knowledge doesn't exist, this can be paralyzing – multiple perspectives co-exist and rarely converge through theory and discourse alone. Robert Lempert [4] frames this as a driver of the political mistrust seen across Western nations: in a simpler past, expert analysis translated into policy with a reasonable chance of hitting its forecasted mark. Today, expert models either sacrifice precision, which the public experiences as "the model was just wrong"; or sacrifice applicability, which experts experience as "the world didn't behave as modelled." Either way, trust erodes.

Robust Decision Making (RDM) aims to break this pattern by shifting the focus from assumptions to decisions – the guiding principle being: *"agree on decisions, not on* *assumptions"* [5]. Rather than seeking an optimal solution under a specific set of forecasts, RDM identifies strategies that perform well across the widest possible range of futures.

The process is structured around four components, the XLRM framework. Uncertainties (X), particularly outside the decision-maker's control, are scoped through stakeholder engagement, drawing deliberately on diverse perspectives to define the broadest conceivable range each uncertain parameter might take. Levers (L) capture what can actually be influenced, again framed by what's possible rather than what's expected or desirable. Metrics (M) are defined by those actually impacted (not necessarily analysts) as thresholds of acceptable outcome rather than targets to optimise for. Finally, relationships (R) connect these components: this is where the model lives, whether built fresh from stakeholder input or, crucially, where existing expert models find their proper place – as one input among many, rather than the arbiter of the decision.

With the framework in place, the model is run across all combinations of uncertainties and levers within the scoped space. These are sampled with equal probability, not weighted by likelihood, making the analysis explicitly possibilistic: we are asking what could happen, not what we think will. The output is a full map of the decision space across all scenarios we consider conceivable.

This map is where RDM earns its value. By filtering for scenarios in which the defined metrics are met, and examining what those scenarios have in common, decision-makers can identify which lever settings are most critical to success; and crucially, the overall trend across thousands of scenarios is far more informative than any single forecast. With levers set to a reasonable position, the share of successful scenarios among all simulated ones gives a direct, intuitive read on robustness: how well does this decision hold up against what we don't know? And because the outputs are visual and scenario-based, they create a genuine basis for discussion – stakeholders can see directly how changes in assumed ranges play out, challenge what's in the model, and refine it iteratively.

RDM is not a rejection of modelling. It is a reassignment of its role – from oracle to instrument. In complex systems, the goal is not to predict the future with precision, but to make decisions that remain sound across a wide range of futures we cannot predict. That is a more honest ambition, and in an increasingly uncertain world, a more useful one.

[1] W. G. Bennis and B. Nanus, Leaders: The Strategies for Taking Charge. New York: Harper & Row, 1986.

[2] J. D. Farmer, *Making Sense of Chaos: A Better Economics for a Better World*. London: Penguin, 2024.

[3] N. N. Taleb, *The Black Swan: The Impact of the Highly Improbable*. New York: Random House, 2007.

[4] R. J. Lempert, "Decision Making under Deep Uncertainty and the Great Acceleration: The Role of Experts and Policy Analysis in a World in Transition," RAND Corporation, Santa Monica, CA, Rep. PE-A3789-1, Apr. 2025. [Online]. Available: https://www.rand.org/pubs/perspectives/PEA3789-1.html

[5] R. J. Lempert, S. W. Popper, and S. C. Bankes, *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*, RAND Corporation, Santa Monica, CA, Rep. MR-1626-RPC, 2003.