

Bounded Rationality with Subjective Evaluations in Enlivened but Truncated Decision Trees¹

Peter J. Hammond: p.j.hammond@warwick.ac.uk

CRETA (Centre for Research in Economic Theory and Applications)
University of Warwick, Coventry CV4 7AL, UK.

2025 February 26th, typeset from ratEnlivenCRETAr.tex

Abstract: A decision-maker is usually assumed to be Bayesian rational, or to maximize subjective expected utility, within a complete and correctly specified decision model. Following the discussion in Hammond (2007) of Schumpeter's (1911, 1934) concept of entrepreneurship, as well as Shackle's (1953) concept of potential surprise, we consider enlivened decision trees whose growth over time cannot be accurately modelled in full detail. An enlivened decision tree involves more severe limitations than a mis-specified model, unforeseen contingencies, or unawareness, all of which are typically modelled with reference to a universal state space large enough to encompass any decision model that an agent may consider. We consider three motivating examples based on: (i) Homer's classic tale of Odysseus and the Sirens; (ii) a two-period linear-quadratic model of portfolio choice; (iii) the game of Chess. Though our novel framework transcends standard notions of risk or uncertainty, for finite decision trees that may be truncated because of bounded rationality, an extended form of Bayesian rationality is still possible, with real-valued subjective evaluations instead of consequences attached to some terminal nodes. Moreover, these subjective evaluations underlie, for example, the kind of Monte Carlo tree search algorithm used by recent chess-playing software packages. [196 words]

Keywords: Bounded Bayesian rationality, consequentialist decision theory, Schumpeterian entrepreneurship, Shackle's potential surprise, truncated decision trees, enlivened decision trees, subjective evaluation of continuation subtrees, Monte Carlo tree search.

JEL Classification: D11, D63, D81, D91.

¹This is a significantly revised version of CRETA working paper # 89, which bore the somewhat less informative title "Bayesian Rationality with Subjective Evaluations in Enlivened Decision Trees".

Prologue

Grau, teurer Freund, ist alle Theorie;
Und grün des Lebens gold'ner Baum.²
— Mephistopheles in Goethe's *Faust*, Part I.³

... he said that to finish [the] poem he could not get along without the house because down in the cellar there was an Aleph. He explained that an Aleph is one of the points in space that contains all other points.

The Aleph's diameter was probably little more than an inch, but all space was there, actual and undiminished. Each thing (a mirror's face, let us say) was infinite things, since I distinctly saw it from every angle of the universe. I saw the Aleph from every point and angle, and in the Aleph I saw the earth and in the earth the Aleph and in the Aleph the earth; I saw my own face and my own bowels; I saw your face; and I felt dizzy and wept, for my eyes had seen that secret and conjectured object whose name is common to all men but which no man has looked upon — the unimaginable universe. I felt infinite wonder, infinite pity. ... for Cantor's Mengenlehre,⁴ [Aleph, or \aleph] is the symbol of transfinite numbers, of which any part is as great as the whole.

Out on the street, going down the stairways inside Constitution Station, riding the subway, every one of the faces seemed familiar to me. I was afraid that not a single thing on earth would ever again surprise me; I was afraid I would never again be free of all I had seen. Happily, after a few sleepless nights, I was visited once more by oblivion.

— Excerpts from Jorge Luis Borges *El Aleph* (1945), translated by Norman Thomas Di Giovanni in collaboration with the author.

²One possible translation is: "Grey, dear friend, is all theory; and green the golden tree of life."

³The subject of this paper provided the content for my last seminar at Stanford before retiring in early 2007. A day or two beforehand, Kenneth Arrow left me a phone message asking if I had been inspired by this quotation from Goethe. While my answer had to be negative, I was left feeling that this should have been the source of my inspiration.

⁴"Mengenlehre" is "set theory" in German.

1 Background and Outline

1.1 Justifying Bayesian Rationality

In decision theory, Bayesian rationality is the hypothesis that a decision-making agent makes choices whose consequences, which are generally lotteries with both risky and uncertain outcomes, maximize the expected value of a Bernoulli utility function $Y \ni y \mapsto u(y) \in \mathbb{R}$ defined on a specified non-empty consequence domain Y . For risky consequences which emerge from what Anscombe and Aumann (1963) describe as a “roulette lottery”, there is by definition an “objective” or hypothetical probability $\pi(\omega) \in [0, 1]$ of each lottery outcome ω in a non-empty finite sample space Ω . For uncertain consequences which emerge from what Anscombe and Aumann (1963) describe as a “horse lottery”, Bayesian rationality requires there to be a “subjective” or personal probability $p(s) \in [0, 1]$ of each lottery outcome or state s in a non-empty finite state space S . A general “Anscombe–Aumann” lottery specifies, for each possible outcome $s \in S$ of a horse lottery with subjective probabilities, a suitable roulette lottery with objective probabilities $\lambda_s(y)$ over consequences y in the non-empty consequence domain Y . Then the appropriate expected utility maximand is the double sum $\sum_{s \in S} p(s) \sum_{y \in Y} \lambda_s(y) u(y)$ involving products of both objective and subjective probabilities.

Past work has offered normative justifications for Bayesian rational behaviour in decision trees based upon:

1. either the “consequentialist” hypothesis set out in Hammond (1988a, b; 1998a, b; 1999) requiring that the range of possible Anscombe–Aumann consequence lotteries which result from prescribed behaviour in any finite decision tree, including any continuation decision tree, should be explicable as the value of a suitably defined choice function defined on the relevant domain of non-empty finite feasible sets of consequence lotteries;
2. or, more recently, the associated concept of prerationality (Hammond, 2022) applied to weak base preference relations \succsim defined on lottery consequence domains.

When either consequentialism or prerationality is assumed, justifying Bayesian rationality does require one additional well-known continuity axiom. This axiom applies to preferences over each “Marschak (1950) triangle” which,

given a triple $\{\lambda, \mu, \nu\}$ of roulette lotteries of which no two are indifferent, is defined as the set $\Delta(\{\lambda, \mu, \nu\})$ of all probability mixtures $q_\lambda\lambda + q_\mu\mu + q_\nu\nu$ of the three lotteries, where the three probability weights q_λ, q_μ, q_ν are all non-negative with $q_\lambda + q_\mu + q_\nu = 1$. After suitable relabelling, one can assume that the strict preference relation \succ satisfies $\lambda \succ \mu$, $\lambda \succ \nu$, and $\mu \succ \nu$. Then, given the corresponding weak preference relation \succeq , the continuity axiom requires that the two sets

$$\{\alpha \in [0, 1] \mid \alpha\lambda + (1 - \alpha)\nu \succeq \mu\} \quad \text{and} \quad \{\alpha \in [0, 1] \mid \mu \succeq \alpha\lambda + (1 - \alpha)\nu\}$$

should both be closed subsets of the unit interval $[0, 1] \subset \mathbb{R}$. Dropping continuity would allow some kind of lexicographic preference relation over lotteries which is not Bayesian rational.

Some of these earlier papers justifying Bayesian rationality also invoked the assumption of dynamic consistency. This assumption requires intended or planned behaviour at the later decision nodes of a tree T to match actual behaviour. Yet in reality actual behaviour is determined only at a decision node of T that happens to be the initial node n_0 of T . At any decision node n of tree T that is not the initial node n_0 of T , behaviour is specified by treating n as the initial node of the *continuation subtree* $T_{\geq n}$ that results from eliminating all the nodes of tree T which do not weakly succeed n . Then, at decision nodes n which come strictly later than the initial node n_0 of T , by specifying actual behaviour without reference to previous intentions or plans, dynamic inconsistency between actions and plans or intentions is entirely ruled out. In this way, dynamic consistency is satisfied by construction.

1.2 Bounded Rationality? Or Bounded Modelling?

“All models are wrong, but some are useful.”

— George Box (1919–2013)

Human ingenuity has led at least some of us to create puzzles and other decision problems in order to amuse or instruct each other. Many children, and some adults, derive satisfaction from solving jigsaw puzzles, or from learning how not to lose at noughts and crosses, otherwise known as tic-tac-toe. Other people try crossword puzzles, or sudoku, or Rubik’s cube. Generations of students take courses in mathematics during which they are expected to learn by solving, or understanding the solutions to, progressively

more demanding exercises. In each of these examples the challenge is to find a perfect solution to a well defined decision problem.

Typical decision problems, however, are not like puzzles or mathematical exercises. Indeed, they are very often far too challenging for full Bayesian rationality to be possible. This recognition, of course, was a key motivation for Simon (1955, 1957) to introduce his concepts of bounded rationality and satisficing. Yet satisficing seems hard to motivate except as the result of some compromise which emerges when the benefits of a more intensive search for a Bayesian optimal decision have been traded off against the additional cost of that search. Thus, satisficing seems to apply better to the choice of what decision model to analyse rather than to the choice of what decision to make within a given model that is being analysed. For this reason, it seems that a more satisfactory fundamental concept may be that of a bounded model.⁵ And in the case of decision trees, an obvious form of bounded model is a truncated tree that results from pruning off one or more entire continuation subtrees.

1.3 Enlivened but Truncated Decision Trees

So, motivated by several examples set out in later sections, including the game of Chess, this paper argues that past work on Bayesian rationality in decision trees is seriously limited in its relevance. This is because of the failure to recognize any possibility that a decision maker’s decision tree may be subject to “enlivenment” in the sense of enriching revisions that are needed in order to recognize possibilities which, though in principle they should have formed part of the original tree, had to be excluded because of computational or other practical modelling limitations.

In order to allow the decision tree to change, even unpredictably, a framework with “enlivened” decision trees is proposed. An entirely myopic agent who follows the old adage “Don’t cross your bridges before you come to them” — which Savage (1963, p. 16) in particular has discussed — will act as though this enlivening is totally irrelevant. This leads to the agent lurching from one model to the next, displaying hubris throughout.

Of course many future decisions and their uncertain consequences cannot be modelled in any detail. Nevertheless, an agent with even a little sophisti-

⁵I should not claim any originality for this thought. See, for example, the discussion of Simon, Shackle, and Chess in Section 8.2 of Earl and Littleboy (2014).

cation should recognize that what matters for any one decision is the current expectation of what, when viewed in retrospect, its ultimate *ex post* value will be. Following ideas that Koopmans (1964) and Kreps (1990, 1992) developed in order to discuss the preference for flexibility, an agent should seek to determine these expected valuations as reliably as possible, using whatever limited evidence is deemed to be relevant, as well as what can be handled within whatever bounded resources the agent can afford to allocate to decision analysis. See also Dekel *et al.* (2001, 2005) and many successors for related ideas in the context of decision making with unforeseen contingencies whose possibility is, nevertheless, foreseen by an apparently omniscient and hubristic decision analyst. See also the work on unawareness in decisions and games by, inter alia, Schipper (2014a, b), Halpern and Rêgo (2014), Grant *et al.* (2015a, b), and especially the related work on growing awareness and reverse Bayesianism by Vierø (2009, 2021) and by Karni and Vierø (2013, 2015, 2017).⁶

We emphasize that the present paper differs from this earlier work on unforeseen contingencies or unawareness by not relying on the existence of any “augmented conceivable state space” of the kind defined in Karni and Vierø (2017, p. 304). Instead, initially we allow the relevant state space to grow entirely unpredictably as a result of the dynamic process that we call “enlivenment”. Specifically, though a decision-making agent may be aware of the possibility of their own unawareness, they are unable even to formulate a practical model which is based, as usual, on a comprehensive space of all conceivably possible states. This enrichment of the previous concepts of unforeseen contingencies or unawareness, which was introduced informally in Hammond (2007), is inspired in part by Schumpeter’s (1911, 1934) concept of entrepreneurship, as well as by Shackle’s (1953) concept of potential surprise.⁷ The concept of an enlivened decision tree was motivated in part by the classical example of Odysseus and the Sirens discussed in Section 3.

That said, a completely specified enlivened decision tree, which is never subject to any further enlivenment, could be regarded as falling within a universal augmented conceivable state space. As discussed above, the existence of such a universal state space raises conceptual problems. To avoid

⁶Other relevant work on unawareness includes the papers published in the special issue of *Mathematical Social Sciences* edited by Schipper (2014a), as well as those cited in Vierø (2021).

⁷See Metcalfe *et al.* (2021) for a recent discussion of links between these two concepts.

these, we recognize the relevance of recursively enlivened decision trees whose growth can never be fully described in a single universal model. Nevertheless, even a fully enlivened tree can still be reduced to a minimally enlivened tree with random outcomes that, instead of consequence lotteries, are subjective evaluations attached to the terminal nodes of a truncated decision tree. As discussed in Section 7.5, this approach to valuing a continuation subtree which can never be completely modelled was the basis of the successful Deep Fritz and then Stockfish open source engines for computer chess. Eventually, however, Stockfish has been supplanted by AlphaZero which results from a special kind of artificial intelligence.

1.4 Outline

Section 2 briefly reviews some distinctions between unbounded and bounded rationality, including as prominent examples of the latter both Simon’s concept of procedural rationality, as well as Manzini and Mariotti’s (2007) “rational shortlist” method.

Next, Section 3 revisits the well known Homeric example of Odysseus and the Sirens. Previous work such as Strotz (1956), Pollak (1968), Hammond (1976) and Elster (1979) has typically regarded this as a prominent example of changing tastes, illustrating the distinction between naïve and sophisticated choice, as well as the potential value of commitment devices. Here, by contrast, this Homeric example is viewed as a mythical decision tree which the sorceress Kirke (or Circe) enlivened as the sage advice that she was offering Odysseus progressed through several stages.⁸

The next two Sections 4 and 5 focus on two more particular examples. The first is of a consumer who, as an investor, chooses a portfolio of financial assets in order to maximize a two-period quadratic utility function subject to a linear budget constraint. Enlivening this consumer’s decision problem could merely affect parameter values, but it could also allow the possibility that new commodities, which may even not yet have been invented, could become relevant.

In Section 5, the second of these two examples concerns the game of Chess, whether played by computers or by humans. It offers a cursory explanation of how Monte Carlo tree simulation can allow a computer algorithm to evaluate positions that arise after possible future moves have been analysed in detail as

⁸In Homeric Greek, the spelling of her name is *Κίρκη*.

far as possible. It also presents a brief case study of a particularly unfortunate human move, described at the time as “the blunder of the century”. This is seen as one particularly prominent player’s failure to revise his bounded model of how the game was likely to proceed.

The main idea of the paper is set out and developed in Sections 6–9. First, Section 6 provides a summary of the key concepts we need to describe Bayesian rational behaviour in classical “unenlivened” finite decision trees. A key tool used in later analysis is the evaluation $v(T)$ of any decision tree T . This is defined as the normalized expected utility generated by any consequence lottery that can result from deciding optimally at each decision node of T . Following the principle of optimality in dynamic programming, this evaluation can be calculated by backward recursion, starting at the terminal nodes of T , each of which has a specified consequence that generally takes the form of a lottery.

Any chess player who was unboundedly rational, and could always make the calculations needed to evaluate any decision tree perfectly, would always play a perfect game. Evidently human players, and even the strongest chess playing computer algorithms, can analyse only bounded decision models. Recognizing this fact, Section 7 introduces the notion of a bounded model in the form of a “truncated” decision tree. With consequences specified only for terminal nodes of a tree that has not been truncated, the backward recursion method of evaluation fails for nodes of a truncated decision tree that may eventually be succeeded by “truncation” nodes rather than by terminal nodes of the original complete tree. Section 7 proposes filling this gap by assigning to each truncation node a real-valued estimate of what the evaluation of the missing continuation subtree starting there would have been if a complete analysis of that subtree were possible.

Section 8 turns at last to a more formal analysis of enlivened decision trees. Initially the focus of is on a special kind of “minimal” enlivenment that is like grafting a new continuation subtree onto the middle of just one edge of the original tree. The initial node of this new enlivening continuation subtree is then an extra node on this edge. This initial extra node is an event node where a horse lottery determines whether or not the decision maker gets taken outside the original tree. There is also discussion of: (i) “simply” enlivened decision trees which result from adding a finite set of minimal enlivenments to the edges of the original tree, at most one to each edge; (ii) “recursively” enlivened decision trees which result when a simply enlivened decision tree is further enlivened.

An agent whose decisions in an enlivened decision tree are fully Bayesian rational is effectively acting as an unboundedly rational agent would if the enlivened decision tree really were the true and complete model of their decision problem. This complete model plays the role of the “augmented conceivable state space” considered by Karni and Vierø (2017, p. 304), amongst others, whose use was criticised in Section 1.3. Such a space seems close in spirit to the fictional device that Borges calls the “Aleph”. Instead, Section 9 weakens Bayesian rationality when facing an enlivened decision tree to the much less demanding requirement of Bayesian rationality in a truncation of the enlivened decision tree, as discussed in Section 7.

In particular, Section 9.3 then states the main result of the paper claiming that the previous characterizations of Bayesian rationality in finite decision trees which were described in Section 1.1 can be adapted to our new setting of enlivened but truncated decision trees. The subsequent Section 9.4 compares the arbitrariness of utilities and subjective probabilities in our model of Bayesian rationality with enlivened evaluations to the arbitrariness of those concepts in the Anscombe and Aumann (1963) model of subjective probability.

The concluding Section 10 starts, in Section 10.1, by a brief discussion of how to extend the results of this paper to the framework in Hammond and Troccoli Moretti (2025) where decision trees may have non-terminal timed consequence nodes. These include consequence nodes with “menu consequences” which depend on the continuation subtree whose initial node is the consequence node. Section 10.2 then analyses briefly the concept of “reverse Bayesianism” as was mentioned in Section 1.3. Next, Section 10.3 offers a brief discussion of recent work by Ullmann-Margalit (2006), Paul (2014, 2015a, b, c) and other philosophers who have introduced the concept of a “transformative experience”. Finally, Section 10.4 offers a brief concluding summary.

2 Beyond Unbounded Rationality

To see a World in a Grain of Sand
And a Heaven in a Wild Flower,
Hold Infinity in the palm of your hand
And Eternity in an hour.
— From William Blake’s “Auguries of Innocence”

2.1 Unbounded versus Procedural Rationality

Simon's (1955, 1957) famous concept of "bounded rationality" may perhaps best be defined by its negation. Decision agents who are *unboundedly rational* make perfect decisions based on perfect models of all the possible acts they could choose, along with all their potential consequences. The result could be the rather disturbing kind of complete model so artfully described in Jorge Luis Borges' short story "El Aleph", from which extracts are quoted in the prologue.

The definition of unbounded rationality in perfect models remains the same no matter whether the consequences are certain (determinate), or else, using the terminology due to Anscombe and Aumann (1963): (i) risky, with hypothetical "objective" probabilities as in a roulette lottery; (ii) uncertain, with personal or "subjective" probabilities as in a horse lottery. Such unbounded rationality would threaten to make games as complicated and enthralling as chess or Go no more interesting than the children's game of noughts and crosses, also known as "tic-tac-toe".⁹ And there would be no such thing as the "law of unintended consequences"; every possible consequence should be calculated, making it in some sense intentional, even as the perhaps unfortunate outcome of a risky decision.

In addition to bounded rationality, Simon advanced the important related idea of "procedural rationality". This recognizes that decision *procedures* could be rational, even if they lead to decisions that are irrational in the sense of violating unbounded rationality. He emphasized concepts like *aspiration level*, along with *satisficing*. The latter appears to mean finding a decision that reaches the aspiration level, and making a decision that seems good enough rather than optimal. But optimal (or even just flexible) search suggests that if the aspiration level is reached quickly and easily, it is too undemanding and so should be raised.

⁹Note that in Chess, the "Lomonosov tablebases" that are distributed online at <http://tb7.chessok.com/> currently specify perfect play starting from any legally possible position *provided* that there is a total of no more than seven pieces of either colour left on the board, including both kings. The usual game of Chess starts, of course, with each of the two players having 16 pieces on the board.

2.2 Rational Shortlists

The normative framework we propose, by contrast, suggests that satisficing behaviour should occur, not within a given decision model, but in choosing how much detail to include in the model and how much to exclude from it. Then ultimately behaviour should be optimal relative to whatever bounded model has been selected for analysis.

One way to use a bounded model involves the “rational shortlist” method that Manzini and Mariotti (2007) introduced to discuss their concept of “sequential rationalizability” for the case of decision problems whose acts have only determinate consequences. Their idea is that, given an inconveniently large feasible set of options, at an initial stage the agent could shortlist a relatively small subset for later serious consideration. Moreover, this shortlist should be small enough to make finding a fully optimal decision amongst those that are shortlisted a manageable decision problem. Thus, any shortlist can be thought of as a bounded model of the feasible set. Also, when it is recognized that observation and/or computation can be costly, work on “rational inattention” inspired by Sims (2003, 2011) and by Hansen and Sargent (2007) considers what bounded decision model may be optimal.

The choice of shortlist can be supposed to emerge rather arbitrarily, even randomly, from some kind of boundedly rational search procedure. Of course, some options may be much more likely to be shortlisted than others. Also, if the composition of the shortlist is regarded as random, the different random variables indicating whether each option belongs to the shortlist may well be correlated.

Once the shortlist has been determined at the first stage, however, it is entirely reasonable to assume that, at a subsequent second stage, the agent indeed selects an optimal element among those that have been shortlisted. That is, choice from within the shortlist satisfies what Simon (1955, 1957) would call “substantive rationality”.

2.3 Other Bounded Decision Models

Shortlisting can be viewed as a particular form of procedural rationality, involving a two-stage procedure. The main point to be made here, however, is that whatever the shortlist may be, it represents a *bounded* model of the full decision problem. Indeed, limitations like the inability of computers to play chess perfectly apply to all difficult decision problems, including most

of those that arise in life rather than in the oversimplified models that are typically analysed and applied by economists and other decision scientists. For this reason, any model we use to inform our decision-making should be flexible enough to allow graceful adaptation to potential changes that any practical model must otherwise ignore.

Suppose an effort really is made to take Simon’s “procedural rationality” idea as seriously as possible. Specifically, it is presumably interesting to explore the implications of assuming that:

1. agents’ time, attention, and computational resources are far too limited for all but simplified models;
2. and in fact they confine themselves to bounded models which are sufficiently simple that they really can find the decision that is optimal within the confines of their bounded model.

Once one recognizes, however, that the model which an agent uses for making decisions is bounded, then one must also recognize that events may eventually force consideration of an expanded or “enlivened” model that includes unmodelled changes.

3 Odysseus and the Sirens Revisited

3.1 A Naïve Sailor’s Model

As our first “classical” example of an enlivened decision tree, we reconsider the Homeric myth of Odysseus and the Sirens. According to this epic myth, naïve sailors whose shortest sea route passed near the Sirens’ island had perhaps in the past used a bounded model of their decision tree like the one illustrated in Figure 1.¹⁰ Specifically, these naïve sailors acted as though they thought that their choice was between:

- either **going near** the Sirens’ island and reaching their destination **early** by a direct route;

¹⁰Note that in Figures 1–4, each square box indicates a terminal node with a consequence described by the word in the box. Apart from these terminal nodes, all other nodes of the tree are decision nodes. Apart from the initial decision node labelled “start”, each other decision node is labelled with a word or phrase describing what is the last action that would lead to that node.

- or **avoiding** the Sirens’ island and arriving **late** after a detour.




Figure 1: Naïve Sailor




Figure 2: Sophisticated Sailor

3.2 A Sophisticated Sailor’s Model

According to Homer, however, Odysseus has the sorceress Kirke as a supernaturally well-informed adviser. She warned Odysseus that the Sirens’ singing had the power to lure unwary sailors to their deaths, and that the meadows on the Sirens’ island were littered with sailors’ bones. So if any naïve sailor came within earshot of the Sirens by choosing **go near** in the decision tree shown in Figure 1, they would find themselves facing instead the decision node marked **go near** in the enlivened and so expanded decision tree shown in Figure 2. At this node in the enlivened tree, their apparent choice would be:

- either **go on** home after hearing the Sirens,
- or **tarry**, enchanted by their singing, and **die** on their island, before ever reaching their intended destination.

Of course, the added feature was that, after hearing the Sirens, no sailor had ever exercised enough will-power to escape the island. This is the essential characteristic of what, in Hammond (1976), was called “potential addict” example of changing tastes. Faced with the decision tree of Figure 2, a sophisticated sailor who understands the persuasive power of the Sirens’ singing would avoid their island and stay out of earshot, even at the cost of only reaching their intended destination after a significant delay.

3.3 Kirke's First Enlivened Model for Odysseus




Figure 3: Kirke's First Model




Figure 4: Kirke's Final Model

Kirke's advice was not confined to a warning, however. Rather routine and unheroic stories about avoiding the Sirens' island and getting back to Ithaca somewhat late by a roundabout route do not constitute memorable epics. Instead Kirke drew attention to the possibility of sailing safely past the Sirens' island, provided the precaution was taken of stopping up all the sailors' ears with wax. Thus, after deciding to approach the Sirens' island, but before getting within earshot, the choice at the node `go near` in Figure 3 would be:

- either `wax` all the crew's ears (including those of Odysseus himself), so none of them hears the Sirens;
- or use `no wax`, like earlier naïve sailors whose bones now litter the Sirens' meadow.

3.4 Kirke's Final Enlivened Model for Odysseus

A much more interesting epic, however, is the one that Homer has given us. Homer had Kirke advise Odysseus on an even better course of action which allowed Odysseus, at least, to hear the Sirens and yet escape with his life.

Indeed, Odysseus was advised that, in addition to arranging for the ears of all his crew to be waxed, he should have himself bound tightly to the mast. Also, his crew should be given strict instructions that, in response to any pleas for release that they see Odysseus making, not only should these pleas be ignored, but also the tightness of his bounds should be increased even more. Thus, Kirke’s final model for Odysseus includes an extra choice node marked **wax** in Figure 4, where the choice is:

- either **bind** Odysseus to the mast, with ears unwaxed so he can **hear** the Sirens,
- or leave Odysseus **free** but with ears waxed like the rest of the crew.

3.5 Toward Enlivened Decision Trees

The earlier naïve sailors whose bones littered the Sirens’ meadow had a model like that in Figure 1. Once they had heard the Sirens’ singing and so learned of their existence, they may have realized too late that a more appropriate model would have been like that in Figure 2. Odysseus (and his crew) were fortunate enough to be provided with a much more useful model, going even beyond Figure 2 to Figure 3 in the first instance, then ultimately to Figure 4.

Each decision tree in Figures 1–4 is lifeless when considered in isolation. The four trees together, however, tell an epic tale of learning. But it is *not* the usual statistical model of learning more and more about the state of the world within a fixed sample space. Rather, the set of possibilities is expanding, as more and more possibilities are included in the enriched model. By introducing the term “enlivened tree”, I have not resisted the temptation to draw an analogy with a live growing tree. Nor of suggesting a strong analogy to the works of Schumpeter (1911, 1934) on innovation, and of Shackle (1953) on “potential surprise” — see Hammond (2007) for further discussion.

4 A Linear–Quadratic Portfolio Problem

4.1 A Two-Period Portfolio Problem

Our first example concerns a consumer with a two-period Bernoulli utility function that takes the quadratic form

$$u(\mathbf{x}_1, \mathbf{x}_2) = -\frac{1}{2}(\mathbf{x}_1 - \mathbf{a}_1)^\top \mathbf{Q}_1(\mathbf{x}_1 - \mathbf{a}_1) - \frac{1}{2}(\mathbf{x}_2 - \mathbf{a}_2)^\top \mathbf{Q}_2(\mathbf{x}_2 - \mathbf{a}_2) \quad (1)$$

Here \mathbf{x}_1 and \mathbf{x}_2 denote finite-dimensional consumption vectors in the two periods, which may possibly have different dimensions, whereas \mathbf{a}_1 and \mathbf{a}_2 are corresponding parameter vectors. Furthermore, assume that \mathbf{Q}_1 , \mathbf{Q}_2 are symmetric and positive definite square matrices of appropriate dimension.

Suppose that the consumer faces two budget constraints, one each period, which can be written as

$$\mathbf{p}_1^\top \mathbf{x}_1 + \mathbf{q}^\top \mathbf{b} = m_1 \quad \text{and} \quad \mathbf{p}_2^\top \mathbf{x}_2 = m_2 + \mathbf{r}^\top \mathbf{b} \quad (2)$$

where \mathbf{b} denotes a finite-dimensional portfolio vector of net asset holdings at the end of period 1, with \mathbf{q} as the asset price vector in period 1, and then \mathbf{r} as the gross return vector. Of course \mathbf{p}_1 and \mathbf{p}_2 denote commodity price vectors each period, both assumed to be strictly positive, whereas $m_1, m_2 \in \mathbb{R}$ are outside wealth transfers. We allow \mathbf{a}_2 , \mathbf{r} and m_2 all to be uncertain, but treat \mathbf{p}_2 as certain, just as Hicks (1946) did when he used point expectations of future prices in his theory of temporary equilibrium.¹¹ For simplicity we also assume that the symmetric matrix \mathbf{Q}_2 is known in period 1. Finally, we assume that the random gross return vector \mathbf{r} is stochastically independent of both random variables \mathbf{a}_2 and m_2 .

4.2 The Second-Period Optimum

By the start of period 2, we assume that the parameter vector \mathbf{a}_2 , the gross return vector \mathbf{r} , and unearned income m_2 have all become known, along with the portfolio vector \mathbf{b} which is pre-determined by the consumer's own choice in period 1. Accordingly, the consumer's second-period optimization problem, which is independent of whatever \mathbf{x}_1 is chosen in period 1, reduces to

$$\max_{\mathbf{x}_2} \left\{ -\frac{1}{2}(\mathbf{x}_2 - \mathbf{a}_2)^\top \mathbf{Q}_2(\mathbf{x}_2 - \mathbf{a}_2) \right\} \quad \text{subject to} \quad \mathbf{p}_2^\top \mathbf{x}_2 = m_2 + \mathbf{r}^\top \mathbf{b} \quad (3)$$

¹¹For a somewhat similar idea, see Myerson (1983).

To solve this constrained maximization problem, introduce the Lagrangian

$$\mathcal{L}_{\lambda_2}(\mathbf{x}_2) = -\frac{1}{2}(\mathbf{x}_2 - \mathbf{a}_2)^\top \mathbf{Q}_2(\mathbf{x}_2 - \mathbf{a}_2) - \lambda_2(\mathbf{p}_2^\top \mathbf{x}_2 - m_2 - \mathbf{r}^\top \mathbf{b}) \quad (4)$$

Then $\mathcal{L}_{\lambda_2}(\mathbf{x}_2)$ is concave as a function of \mathbf{x}_2 . So it is maximized at any point \mathbf{x}_2 that satisfies the first-order condition

$$\mathbf{0} = \mathcal{L}'_{\lambda_2}(\mathbf{x}_2) = -(\mathbf{x}_2 - \mathbf{a}_2)^\top \mathbf{Q}_2 - \lambda_2 \mathbf{p}_2^\top \quad (5)$$

Because \mathbf{Q}_2 is assumed to be positive definite and so invertible, this first-order condition is evidently equivalent to

$$(\mathbf{x}_2 - \mathbf{a}_2)^\top = -\lambda_2 \mathbf{p}_2^\top \mathbf{Q}_2^{-1} \quad (6)$$

or, after transposing and rearranging, to

$$\mathbf{x}_2 = \mathbf{a}_2 - \lambda_2 \mathbf{Q}_2^{-1} \mathbf{p}_2 \quad (7)$$

Substituting this into the budget equation in (3) gives

$$\mathbf{p}_2^\top \mathbf{x}_2 = \mathbf{p}_2^\top (\mathbf{a}_2 - \lambda_2 \mathbf{Q}_2^{-1} \mathbf{p}_2) = m_2 + \mathbf{r}^\top \mathbf{b} \quad (8)$$

implying that

$$\lambda_2 = \frac{\mathbf{p}_2^\top \mathbf{a}_2 - m_2 - \mathbf{r}^\top \mathbf{b}}{\mathbf{p}_2^\top \mathbf{Q}_2^{-1} \mathbf{p}_2} \quad (9)$$

Note that the solution λ_2 exists because $\mathbf{p}_2 \neq \mathbf{0}$ and \mathbf{Q}_2 is positive definite. Finally, we can combine (9) with (7) to determine the optimal demand vector, which is

$$\mathbf{x}_2^* = \mathbf{a}_2 - \frac{\mathbf{p}_2^\top \mathbf{a}_2 - m_2 - \mathbf{r}^\top \mathbf{b}}{\mathbf{p}_2^\top \mathbf{Q}_2^{-1} \mathbf{p}_2} \mathbf{Q}_2^{-1} \mathbf{p}_2 \quad (10)$$

Of course, for this solution to be economically sensible, we should require that $\lambda_2 \geq 0$, or equivalently, that $\mathbf{p}_2^\top \mathbf{a}_2 \geq m_2 + \mathbf{r}^\top \mathbf{b}$. Because this inequality involves the asset vector \mathbf{b} chosen in the first period, we will return to this issue later after deriving the consumer's optimal decisions in the first period.

Note that this solution implies that *ex post*, after \mathbf{a}_2 , \mathbf{r} and m_2 have all become known and \mathbf{x}_2^* has been chosen optimally, equations (7) and (9) imply that the consumer's maximized second period utility is

$$\begin{aligned} -\frac{1}{2}(\mathbf{x}_2^* - \mathbf{a}_2)^\top \mathbf{Q}_2(\mathbf{x}_2^* - \mathbf{a}_2) &= -\frac{1}{2} \lambda_2^2 \mathbf{p}_2^\top \mathbf{Q}_2^{-1} \mathbf{Q}_2 \mathbf{Q}_2^{-1} \mathbf{p}_2 \\ &= -\frac{(\mathbf{p}_2^\top \mathbf{a}_2 - m_2 - \mathbf{r}^\top \mathbf{b})^2}{2 \mathbf{p}_2^\top \mathbf{Q}_2^{-1} \mathbf{p}_2} \end{aligned} \quad (11)$$

4.3 First-Period Expected Utility

Coming back to the first period, we have assumed that \mathbf{p}_2 and \mathbf{Q}_2 are both known in advance. So after using (11), the *ex ante* expected value of the intertemporal Bernoulli utility function (1) can be expressed as the function

$$v(\mathbf{x}_1, \mathbf{b}) = -\frac{1}{2}(\mathbf{x}_1 - \mathbf{a}_1)^\top \mathbf{Q}_1(\mathbf{x}_1 - \mathbf{a}_1) - \frac{\mathbb{E}(\mathbf{p}_2^\top \mathbf{a}_2 - m_2 - \mathbf{r}^\top \mathbf{b})^2}{2\mathbf{p}_2^\top \mathbf{Q}_2^{-1} \mathbf{p}_2} \quad (12)$$

of the first-period choice variables \mathbf{x}_1 and \mathbf{b} . The numerator of the fraction in the second term of the right-hand side of (12) can be expanded as

$$\begin{aligned} \mathbb{E}(\mathbf{p}_2^\top \mathbf{a}_2 - m_2 - \mathbf{r}^\top \mathbf{b})^2 \\ = \mathbb{E}(\mathbf{p}_2^\top \mathbf{a}_2 - m_2)^2 - 2\mathbb{E}[(\mathbf{p}_2^\top \mathbf{a}_2 - m_2)(\mathbf{r}^\top \mathbf{b})] + \mathbb{E}(\mathbf{r}^\top \mathbf{b})^2 \end{aligned} \quad (13)$$

Let $\bar{\mathbf{a}}_2 := \mathbb{E}\mathbf{a}_2$, $\bar{m}_2 := \mathbb{E}m_2$ and $\bar{\mathbf{r}} := \mathbb{E}\mathbf{r}$ denote the respective means, all of which are assumed to exist. Our assumption that \mathbf{r} is stochastically independent of \mathbf{a}_2 and m_2 implies that the middle term on the right-hand side of (13) reduces to

$$\mathbb{E}[(\mathbf{p}_2^\top \mathbf{a}_2 - m_2)(\mathbf{r}^\top \mathbf{b})] = (\mathbf{p}_2^\top \bar{\mathbf{a}}_2 - \bar{m}_2)(\bar{\mathbf{r}}^\top \mathbf{b}) \quad (14)$$

As for the last term on the right-hand side of (13), note that

$$(\mathbf{r}^\top \mathbf{b})^2 = (\mathbf{b}^\top \mathbf{r})(\mathbf{r}^\top \mathbf{b}) = \mathbf{b}^\top (\mathbf{r}\mathbf{r}^\top) \mathbf{b} \quad \text{and so} \quad \mathbb{E}(\mathbf{r}^\top \mathbf{b})^2 = \mathbf{b}^\top \mathbf{R} \mathbf{b} \quad (15)$$

where \mathbf{R} denotes the symmetric square matrix $\mathbb{E}[\mathbf{r}\mathbf{r}^\top]$ of second moments of returns, which we also assume exists. The matrix \mathbf{R} is positive definite under the assumption that the second moment $\mathbb{E}(\mathbf{r}^\top \mathbf{b})^2$ of the return to any portfolio $\mathbf{b} \neq \mathbf{0}$ is always positive.

Substituting from (14) and (15) in (13) gives

$$\begin{aligned} \mathbb{E}(\mathbf{p}_2^\top \mathbf{a}_2 - m_2 - \mathbf{r}^\top \mathbf{b})^2 &= \mathbb{E}(\mathbf{p}_2^\top \mathbf{a}_2 - m_2)^2 - 2(\mathbf{p}_2^\top \bar{\mathbf{a}}_2 - \bar{m}_2)\bar{\mathbf{r}}^\top \mathbf{b} + \mathbf{b}^\top \mathbf{R} \mathbf{b} \\ &= c + (\mathbf{b}^* - \mathbf{b})^\top \mathbf{R}(\mathbf{b}^* - \mathbf{b}) \end{aligned} \quad (16)$$

where $\mathbf{b}^{*\top} \mathbf{R} = (\mathbf{p}_2^\top \bar{\mathbf{a}}_2 - \bar{m}_2)\bar{\mathbf{r}}^\top$, implying that $\mathbf{b}^* = \mathbf{R}^{-1}\bar{\mathbf{r}}(\mathbf{p}_2^\top \bar{\mathbf{a}}_2 - \bar{m}_2)$, and also

$$c = \mathbb{E}(\mathbf{p}_2^\top \mathbf{a}_2 - m_2)^2 - \mathbf{b}^{*\top} \mathbf{R} \mathbf{b}^* = \mathbb{E}(\mathbf{p}_2^\top \mathbf{a}_2 - m_2)^2 - (\mathbf{p}_2^\top \bar{\mathbf{a}}_2 - \bar{m}_2)^2 \bar{\mathbf{r}}^\top \mathbf{R}^{-1} \bar{\mathbf{r}} \quad (17)$$

Finally, therefore, after ignoring an irrelevant additive constant, the consumer's first-period maximand can be written as the quadratic form

$$v(\mathbf{x}_1, \mathbf{b}) = -\frac{1}{2}(\mathbf{x}_1 - \mathbf{a}_1)^\top \mathbf{Q}_1(\mathbf{x}_1 - \mathbf{a}_1) - \frac{1}{2}(\mathbf{b}^* - \mathbf{b})^\top \mathbf{S}(\mathbf{b}^* - \mathbf{b}) \quad (18)$$

where $\mathbf{S} := \mathbf{R}/\mathbf{p}_2^\top \mathbf{Q}_2^{-1} \mathbf{p}_2$.

4.4 The First-Period Optimization Problem

The consumer's first-period optimization is therefore to maximise the function (18) w.r.t. \mathbf{x}_1 and \mathbf{b} , subject to the budget constraint $\mathbf{p}_1^\top \mathbf{x}_1 + \mathbf{q}^\top \mathbf{b} = m_1$. We solve this constrained maximization problem by introducing the Lagrangian

$$\begin{aligned} \mathcal{L}_{\lambda_1}(\mathbf{x}_1, \mathbf{b}) = & -\frac{1}{2}(\mathbf{x}_1 - \mathbf{a}_1)^\top \mathbf{Q}_1(\mathbf{x}_1 - \mathbf{a}_1) - \frac{1}{2}(\mathbf{b}^* - \mathbf{b})^\top \mathbf{S}(\mathbf{b}^* - \mathbf{b}) \\ & - \lambda_1(\mathbf{p}_1^\top \mathbf{x}_1 + \mathbf{q}^\top \mathbf{b} - m_1) \end{aligned} \quad (19)$$

which is concave as a function of $(\mathbf{x}_1, \mathbf{b})$, so is maximized w.r.t. $(\mathbf{x}_1, \mathbf{b})$ when the two first-order conditions

$$\begin{aligned} \mathbf{0} &= \mathcal{L}'_{\lambda_1, \mathbf{x}_1} = -(\mathbf{x}_1 - \mathbf{a}_1)^\top \mathbf{Q}_1 - \lambda_1 \mathbf{p}_1^\top \\ \text{and } \mathbf{0} &= \mathcal{L}'_{\lambda_1, \mathbf{b}} = (\mathbf{b}^* - \mathbf{b})^\top \mathbf{S} - \lambda_1 \mathbf{q}^\top \end{aligned} \quad (20)$$

are both satisfied. Because both \mathbf{Q}_1 and \mathbf{S} are positive definite and so invertible, these first-order conditions are equivalent to

$$(\mathbf{x}_1 - \mathbf{a}_1)^\top = -\lambda_1 \mathbf{p}_1^\top \mathbf{Q}_1^{-1} \quad \text{and} \quad (\mathbf{b}^* - \mathbf{b})^\top = \lambda_1 \mathbf{q}^\top \mathbf{S}^{-1} \quad (21)$$

or, after transposing and rearranging, to

$$\mathbf{x}_1 = \mathbf{a}_1 - \lambda_1 \mathbf{Q}_1^{-1} \mathbf{p}_1 \quad \text{and} \quad \mathbf{b} = \mathbf{b}^* - \lambda_1 \mathbf{S}^{-1} \mathbf{q} \quad (22)$$

Substituting these into the budget equation gives

$$\mathbf{p}_1^\top \mathbf{x}_1 + \mathbf{q}^\top \mathbf{b} = \mathbf{p}_1^\top (\mathbf{a}_1 - \lambda_1 \mathbf{Q}_1^{-1} \mathbf{p}_1) + \mathbf{q}^\top (\mathbf{b}^* - \lambda_1 \mathbf{S}^{-1} \mathbf{q}) = m_1 \quad (23)$$

implying that

$$\lambda_1 = \frac{\mathbf{p}_1^\top \mathbf{a}_1 + \mathbf{q}^\top \mathbf{b}^* - m_1}{\mathbf{p}_1^\top \mathbf{Q}_1^{-1} \mathbf{p}_1 + \mathbf{q}^\top \mathbf{S}^{-1} \mathbf{q}} \quad (24)$$

Note that this is well defined because $\mathbf{p}_1 \neq \mathbf{0}$ and $\mathbf{q} \neq \mathbf{0}$, whereas both symmetric matrices \mathbf{Q}_1 and \mathbf{S} are positive definite, and so invertible with inverses that are positive definite.¹² Finally, we can use (22) and (24) in

¹²A standard result in matrix theory is that any symmetric $n \times n$ matrix \mathbf{A} can be diagonalized, in the sense that there exists an orthogonal matrix \mathbf{E} (meaning that $\mathbf{E}^{-1} = \mathbf{E}^\top$) and a diagonal matrix \mathbf{D} such that $\mathbf{E}\mathbf{A}\mathbf{E}^\top = \mathbf{D}$ and so $\mathbf{A} = \mathbf{E}^\top \mathbf{D} \mathbf{E}$. Then it is easy to see that the following four conditions are all logically equivalent: (i) \mathbf{A} is positive definite; (ii) all diagonal elements of \mathbf{D} are positive; (iii) \mathbf{D}^{-1} exists and all its diagonal elements are positive; (iv) $\mathbf{A}^{-1} = \mathbf{E}^\top \mathbf{D}^{-1} \mathbf{E}$ exists and is positive definite.

order to determine the optimal commodity and asset demand vectors, which are

$$\mathbf{x}_1^* = \mathbf{a}_1 - \frac{\mathbf{p}_1^\top \mathbf{a}_1 + \mathbf{q}^\top \mathbf{b}^* - m_1}{\mathbf{p}_1^\top \mathbf{Q}_1^{-1} \mathbf{p}_1 + \mathbf{q}^\top \mathbf{S}^{-1} \mathbf{q}} \mathbf{Q}_1^{-1} \mathbf{p}_1 \quad (25)$$

$$\text{and } \mathbf{b} = \mathbf{b}^* - \frac{\mathbf{p}_1^\top \mathbf{a}_1 + \mathbf{q}^\top \mathbf{b}^* - m_1}{\mathbf{p}_1^\top \mathbf{Q}_1^{-1} \mathbf{p}_1 + \mathbf{q}^\top \mathbf{S}^{-1} \mathbf{q}} \mathbf{S}^{-1} \mathbf{q} \quad (26)$$

Of course, for this solution to be economically sensible, we should require that $\lambda_1 \geq 0$, or equivalently, that

$$\mathbf{p}_1^\top \mathbf{a}_1 + \mathbf{q}^\top \mathbf{b}^* = \mathbf{p}_1^\top \mathbf{a}_1 + \mathbf{q}^\top \mathbf{R}^{-1} \bar{\mathbf{r}} (\mathbf{p}_2^\top \bar{\mathbf{a}}_2 - \bar{m}_2) \geq m_1 \quad (27)$$

Furthermore, for the second-period solution we found previously to be economically sensible, we should require also that $\mathbf{p}_2^\top \mathbf{a}_2 \geq m_2 + \mathbf{r}^\top \mathbf{b}$. Because the two random variables $\mathbf{p}_2^\top \mathbf{a}_2 - m_2$ and $\mathbf{r}^\top \mathbf{b}$ are independent, this requirement implies that there must be a real number α for which, given the optimal choice of \mathbf{b} , one has

$$\mathbf{p}_2^\top \mathbf{a}_2 - m_2 \geq \alpha \geq \mathbf{r}^\top \mathbf{b} \quad (28)$$

for almost all possible values of the random pair $(\mathbf{p}_2^\top \mathbf{a}_2 - m_2, \mathbf{r}^\top \mathbf{b}) \in \mathbb{R}^2$.

4.5 An Enlivened Decision Problem

To enliven this linear-quadratic decision model, we consider the possibility that unforeseeable changes occur after the pair $(\mathbf{x}_1, \mathbf{b})$ has already been chosen in period 1. In general, there could be a new second period objective

$$-\frac{1}{2}(\mathbf{x}_2^+ - \mathbf{a}_2^+)^\top \mathbf{Q}_2^+ (\mathbf{x}_2^+ - \mathbf{a}_2^+) \quad (29)$$

in which the dimension of the vectors \mathbf{x}_2^+ , \mathbf{a}_2^+ and the corresponding dimension of the positive definite square matrix \mathbf{Q}_2^+ may have increased, perhaps because of new commodities. Of course, the second-period budget constraint must also change; we write it as

$$\mathbf{p}_2^{+\top} \mathbf{x}_2^+ \leq \mathbf{r}^\top \mathbf{b} + m_2 \quad (30)$$

with the same asset vector \mathbf{b} as before, since that is already determined by the consumer's decisions in period 1. The joint distribution of the triple $(\mathbf{a}_2^+, m_2, \mathbf{r})$ may also change, as indeed it must if the dimension of \mathbf{a}_2^+ exceeds that of \mathbf{a}_2 .

If these changes could be known in advance, then in period 1 the consumer would face the problem of maximizing, instead of the quadratic evaluation function $v(\mathbf{x}_1, \mathbf{b})$ defined by (18), a revised quadratic objective function

$$v^+(\mathbf{x}_1, \mathbf{b}) = -\frac{1}{2}(\mathbf{x}_1 - \mathbf{a}_1)^\top \mathbf{Q}_1(\mathbf{x}_1 - \mathbf{a}_1) - \frac{1}{2}(\mathbf{b}^{+*} - \mathbf{b})^\top \mathbf{S}^+(\mathbf{b}^{+*} - \mathbf{b}) \quad (31)$$

of the same choice variables \mathbf{x}_1 and \mathbf{b} , subject to the same first-period budget constraint $\mathbf{p}_1^\top \mathbf{x}_1 + \mathbf{q}^\top \mathbf{b} = m_1$ as in (2). What has changed, however, are the vector parameter \mathbf{b}^{+*} and matrix parameter \mathbf{S}^+ which appear in the last term of (31), whose changes are now entirely unpredictable. Enlivenment requires recognizing that these parameters must be treated as themselves uncertain. A Bayesian rational consumer who remains convinced that some quadratic model is still appropriate will, by definition, hold some subjective probability beliefs concerning the unpredictable pair $(\tilde{\mathbf{b}}^*, \tilde{\mathbf{S}})$ of parameters that characterize each member of the parametric family of quadratic evaluation functions

$$\tilde{v}(\mathbf{x}_1, \mathbf{b}; \tilde{\mathbf{b}}^*, \tilde{\mathbf{S}}) \equiv -\frac{1}{2}(\mathbf{x}_1 - \mathbf{a}_1)^\top \mathbf{Q}_1(\mathbf{x}_1 - \mathbf{a}_1) - \frac{1}{2}(\tilde{\mathbf{b}}^* - \mathbf{b})^\top \tilde{\mathbf{S}}(\tilde{\mathbf{b}}^* - \mathbf{b}) \quad (32)$$

Rationality, in the sense of subjective expected utility maximization, requires optimal policy in period 1 to maximize the expected value $\widehat{\mathbb{E}}[\tilde{v}(\mathbf{x}_1, \mathbf{b}; \tilde{\mathbf{b}}^*, \tilde{\mathbf{S}})]$ of the function (32) w.r.t. probabilistic beliefs concerning the parameter pair $(\tilde{\mathbf{b}}^*, \tilde{\mathbf{S}})$. Such an expectation, however, after ignoring an irrelevant additive constant, can be expressed in the convenient form

$$\widehat{\mathbb{E}}[\tilde{v}(\mathbf{x}_1, \mathbf{b}; \tilde{\mathbf{b}}^*, \tilde{\mathbf{S}})] \equiv -\frac{1}{2}(\mathbf{x}_1 - \mathbf{a}_1)^\top \mathbf{Q}_1(\mathbf{x}_1 - \mathbf{a}_1) - \frac{1}{2}(\widehat{\mathbf{b}}^* - \mathbf{b})^\top \widehat{\mathbf{S}}(\widehat{\mathbf{b}}^* - \mathbf{b}) \quad (33)$$

This involves the appropriate subjective expected value $\widehat{\mathbf{S}} := \widehat{\mathbb{E}}[\tilde{\mathbf{S}}]$ of the random matrix $\tilde{\mathbf{S}}$. Note that the matrix $\widehat{\mathbf{S}}$ is positive definite, and so invertible, as the expected value of the random positive definite matrix $\tilde{\mathbf{S}}$.¹³ This allows the vector $\widehat{\mathbf{b}}^*$ to be chosen uniquely so that it satisfies the first-order condition $\widehat{\mathbf{S}} \widehat{\mathbf{b}}^* = \widehat{\mathbb{E}}[\tilde{\mathbf{S}} \tilde{\mathbf{b}}^*]$, implying that

$$\widehat{\mathbf{b}}^* = \widehat{\mathbf{S}}^{-1} \widehat{\mathbb{E}}[\tilde{\mathbf{S}} \tilde{\mathbf{b}}^*] = (\widehat{\mathbb{E}}[\tilde{\mathbf{S}}])^{-1} \widehat{\mathbb{E}}[\tilde{\mathbf{S}} \tilde{\mathbf{b}}^*] \quad (34)$$

¹³To show this, note that if the random symmetric matrix $\tilde{\mathbf{S}}$ is almost surely positive definite, then for all $\mathbf{u} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ one has $\mathbf{u}^\top (\widehat{\mathbb{E}}[\tilde{\mathbf{S}}]) \mathbf{u} = \widehat{\mathbb{E}}(\mathbf{u}^\top \tilde{\mathbf{S}} \mathbf{u}) > 0$.

5 Computer Chess

5.1 Simplified Chess

Consider the decision problem faced by a chess player who has to choose a move when confronted by a known position denoted by n_0 . To specify this position requires saying whose turn it is to move, and what piece, if any, occupies each of the 64 squares on the board.¹⁴ Then let $N_1 := N_{+1}(n_0)$ denote the set consisting of all those positions that can be reached by a move which is legal in position n_0 .

Recall that, in the game of Chess, a player's King is in check just in case it is attacked by an opponent's piece, in the sense that, in the absence of an intervening move, that piece could capture the King. A player's move is legal only if it does not leave that player's King in check. If the player whose turn it is to move has no legal move, then: (i) either that player's King in check, in which case that player has been checkmated and loses the game; (ii) or that player's King is not in check, in which case there is a stalemate and the game is a draw.

Following the famous result of Zermelo (1913), as well as von Neumann's (1928) pioneering analysis of maximin or minimax strategies in two-person "zero-sum" games of perfect information, given best play by both the White and Black players, there is an objective *result function*

$$N_1 \ni n_1 \mapsto r^+(n_1) \in \{W, D, L\} \quad (35)$$

This function maps each possible position $n_1 \in N_1$ to a determinate *result* $r^+(n_1) \in \{W, D, L\}$ of the game that, for the player who is about to move, is either a win (W), or a draw (D), or a loss (L). This result can be converted into a payoff using a scoring rule such as 1 for a win for White, or -1 for a win for Black, but 0 for a draw. Then, given a continuation subgame of Chess that starts from the position n , the result of best play by both players

¹⁴Actually, even in a simplified version of chess — without either clocks that are used to enforce limits on each player's total thinking time, or drawing rules that go beyond stalemate, threefold repetition, or perpetual check — the rules of chess specify that: (i) castling is disallowed if either the king or relevant rook has ever been moved previously; (ii) a pawn can capture an opposition pawn *en passant*, but only immediately after the pawn that is about to be captured has advanced two squares from its initial position. So there are many chess positions whose full description requires significantly more information.

in that subgame will be given by an objective *evaluation function*

$$N_1 \ni n_1 \mapsto v(n_1) \in \{1, 0, -1\} \tag{36}$$

For the player whose turn it is to move at n_0 , a move from n_0 to n_1 is optimal if and only if:

1. n_1 maximizes the evaluation function $v(n_1)$ in case it is White's turn to move at n_0 ;
2. n_1 minimizes the evaluation function $v(n_1)$ in case it is Black's turn to move at n_0 .

The objective normalized valuation function in (36) can only be computed, however, for a few relatively simple positions where:

- either it can be proved that, in a small number of moves, one side can force a win due to checkmate, or else, should they wish, a draw due to either (i) stalemate; (ii) a threefold repetition of the position; (iii) perpetual check;
- or alternatively, there are no more than 7 pieces on the board, including both Kings, in which case the Lomonosov “endgame tablebase” software cited in footnote 7 of Section 2.1 will specify what is the result of the game if both players follow maximin strategies.

Thus, in choosing what move to make at n_0 , and so what should be the next position $n_1 \in N_1$ on the board, a player is typically forced to come up with subjective beliefs regarding the payoff function. These beliefs can be guided by looking ahead a few moves. But unless one can calculate with certainty a way to force a simple position whose evaluation is definitely known, ultimately one has to assign such evaluations to many such positions a few moves ahead. In this way, one constructs a subjective *evaluation function* mapping chess positions into subjectively expected payoffs. Computer chess programs for doing this involve algorithms that are good, even superhuman, but are still necessarily imperfect. Currently some of the most effective software uses an algorithm based on *Monte Carlo tree search* (MCTS), which is further discussed in Section 7.5 — see Browne *et al.* (2012) for a general survey that has been widely cited in the computer science literature. Applied to Chess, in order to evaluate a given position n , MCTS considers many simulated continuation subgames that all start in position n , but then introduces a little

carefully controlled randomness into the routine for choosing each ensuing move. Then the final evaluation of any position n is the average score over all the simulated games that start in position n .

5.2 Real Chess

Real chess is considerably more complicated. For one thing, a player about to move can claim a draw by demonstrating that the next move can be chosen either to repeat the same position a third time, or so that both players will have made at least 50 moves without either a piece being captured or a pawn being moved. Also, the game usually ends with either: (i) one player who is losing choosing to resign; or (ii) with both players agreeing to a draw when both judge that they have an insufficient chance of winning.

Finally, there are time limits monitored by a chess clock, or actually a coupled pair of clocks, one for each player, which displays how much remaining total time that player has available before the next time control. Whenever either player has just made a move, they can press a lever that simultaneously stops their own clock and starts the opponent's. These additional considerations make the description of any chess position n rather more complicated, since it must include, for instance, how much more time each player can use before they would lose on time.

5.3 Human Failure in a Bounded Model

Human chess experts exercise their skill by focusing attention on only a small number of plausible moves in each position. Given any legal chess position n_0 , consider the set $N_1 := N_{+1}(n_0)$ of all possible positions n_1 that can result after a legal move to n_1 is made from the position n_0 . Chess experts discern that many members n_1 of N_1 , though allowable, are too inferior to deserve much, if any, consideration. Of course, human chess experts are also very good at judging the value of any position n_1 that they might think of moving to. In this sense, they have good bounded models.

But, being merely human, even the very best players' models and evaluations of different positions may sometimes be grossly deficient. Witness how in 2006 Vladimir Kramnik, then the world champion, committed the "blunder of the century" by overlooking a checkmate in one move, which led

to an immediate loss. This blunder was during the second game of a match of six games played against the computer program Deep Fritz.¹⁵




Figure 5: Deep Fritz v. Kramnik, Game 2

In this game, Deep Fritz was playing with the White pieces. Its last move before the position shown in Figure 5 was its 34th. The move was 34. Ne6×f8. This notation signifies that White’s knight, which had been on square e6, was used to take the Black piece, actually a rook, which had been on square f8. In response, Kramnik (as Black) blundered horribly by playing the queen move 34 . . . Qa7–e3, as indicated by the arrow in Figure 5, thereby reaching the position shown in that Figure. Whereupon the computer program Deep Fritz promptly indicated that its next move, the queen move 35. Qe4–h7, would win at once by giving checkmate for White.

There must be thousands if not millions of chess players not nearly as strong as Kramnik who, if they were to be shown the position immediately before the move 34 . . . Qa7–e3, would certainly notice that White was threatening to make the move 35. Qe4–h7 checkmate. So how did Kramnik overlook it?

A good clue lies in the following observation that Kramnik himself offered during the press conference that was held minutes after the game ended:

“It was actually not only about the last move. . . . I calculated the line many, many times, rechecking myself. I already calculated this line when I played 29 . . . Qa7, and after each move I was

¹⁵See <https://en.chessbase.com/post/how-could-kramnik-overlook-the-mate->

recalculating, again, and again, and finally I blundered mate in one. Actually it was the first time that it happened to me, and I cannot really find any explanation. I was not feeling tired, I think I was calculating well during the whole game ... It's just very strange, I cannot explain it."

One way to interpret this is that Kramnik as Black, already when choosing his 29th move, had been planning to make what turned out subsequently to be the disastrous move 34 ... Qa7–e3 which is shown in Figure 5. After all, if in Figure 5 the Black king had been on square g8 rather than h8, then Black could have responded to 35. Qe4–h7 check with the move 35. ... K×f8. In the actual game, however, even at the last moment when he was about to make the move 34 ... Qa7–e3, he failed to modify his mental model of the game, as he should have, by recognizing that this move allowed a mate in one move by his opponent. In the end it was as if Kramnik had become so fixated on the earlier plan to choose 34 ... Qa7–e3 that he failed to recognize the possibility of being checkmated immediately after making that fatal move.

6 Bayesian Rationality in Decision Trees

6.1 Roulette Lotteries

Following the terminology of Anscombe and Aumann (1963), given any non-empty set Z , let $\Delta(Z)$ denote the set of all *roulette lotteries* or *simple probability measures*. These take the form of functions $Z \ni z \mapsto \lambda(z) \in [0, 1]$ with a *finite support* $\text{supp } \lambda \subseteq Z$ such that

$$\lambda(z) > 0 \iff z \in \text{supp } \lambda \quad \text{and} \quad \sum_{z \in Z} \lambda(z) = \sum_{z \in \text{supp } \lambda} \lambda(z) = 1 \quad (37)$$

Then, given any $z \in Z$, let $Z \ni z' \mapsto \delta_z(z') \in \Delta(Z)$ denote the unique *degenerate* lottery that satisfies $\delta_z(z) = 1$. Also, whenever Z is a finite set, let $\Delta^0(Z)$ denote the set of *fully supported* lotteries $\lambda \in \Delta(Z)$ that satisfy $\text{supp } \lambda = Z$, or equivalently, $\lambda(z) > 0$ for all $z \in Z$. Finally, given any $\lambda, \mu \in \Delta(Z)$ and any scalar $\alpha \in [0, 1]$, let $\nu := \alpha \lambda + (1 - \alpha) \mu \in \Delta(Z)$ denote the *lottery mixture* $Z \ni z \mapsto \nu(z) \in [0, 1]$ which, for all $z \in Z$, satisfies

$$\nu(z) = [\alpha \lambda + (1 - \alpha) \mu](z) = \alpha \lambda(z) + (1 - \alpha) \mu(z) \quad (38)$$

6.2 Anscombe–Aumann Consequence Lotteries

The hypothesis of Bayesian rationality, or subjective expected utility maximization, applies when there is a non-empty *state space* S of possible states of the world s on which the subjective probability mapping $S \ni s \mapsto p(s) \in [0, 1]$ is defined, where $\sum_{s \in S} p(s) = 1$. Following Anscombe and Aumann (1963) once again, we assume that S is finite. Also, following their terminology which was described in Section 1.1, any random process for determining an uncertain state of the world will be described as a “horse lottery”.

Bayesian rationality concerns preferences over Anscombe–Aumann consequence lotteries. By definition, these may involve both risk, due to roulette lotteries, and uncertainty, due to horse lotteries. Let Y denote a non-empty consequence domain, and then let $\Delta(Y)$ denote the domain of roulette lotteries over Y .

Next, given the finite set S of states s and the consequence domain Y , for each state $s \in S$, let Y_s be a copy of Y .¹⁶ Then let

$$L^S(Y) := \prod_{s \in S} \Delta(Y_s) = \{ \langle \lambda_s \rangle_{s \in S} \mid \forall s \in S : \lambda_s \in \Delta(Y_s) \} \quad (39)$$

denote the space of *Anscombe–Aumann lotteries*, or *AA lotteries*, in the form of lists $\langle \lambda_s \rangle_{s \in S}$ or mappings $S \ni s \mapsto \lambda_s \in \Delta(Y)$. Each such mapping specifies a combination of, first, a horse lottery that determines a state $s \in S$, followed second by a state-dependent roulette lottery λ_s that determines a consequence $y \in Y$.

¹⁶This is the case of a *state-independent consequence domain*, which we assume in order to simplify notation. The more general case of a *state-dependent consequence domain* occurs when Y_s depends on s . In this case, let $Y^\cup := \cup_{s \in S} Y_s$ denote the *union domain* of all consequences y that are feasible in some state $s \in S$. Then there may be a *state-dependent utility function* $D \ni (s, y) \mapsto u(s, y) \rightarrow \mathbb{R}$ defined on the domain $D := \{(s, y) \in S \times Y^\cup \mid y \in Y_s\}$ of feasible state–consequence pairs. Such state-dependent utility functions have been studied in Drèze (1962), Karni (1985), Schervish et al. (1990), Drèze and Rustichini (2004), and Seidenfeld et al. (2010). See Hammond (1998b, 1999, 2022) for a unified treatment which derives a state-independent utility function even in the general case when the consequence domain is state-dependent.

6.3 Choice from Pair Sets and Base Preferences

Let $\mathcal{F}_{\setminus\emptyset}(L^S(Y))$ denote the family of non-empty finite subsets of the AA-lottery domain $L^S(Y)$. A *choice function* on this lottery domain is a mapping

$$\mathcal{F}_{\setminus\emptyset}(L^S(Y)) \ni F \mapsto C(F) \in \mathcal{F}_{\setminus\emptyset}(L^S(Y)) \quad (40)$$

that, for each non-empty *feasible set* $F \in \mathcal{F}_{\setminus\emptyset}(L^S(Y))$, determines a non-empty *choice set* $C(F) \in \mathcal{F}_{\setminus\emptyset}(L^S(Y))$ satisfying $C(F) \subseteq F$.

Corresponding to any choice function $F \mapsto C(F)$ on $\mathcal{F}_{\setminus\emptyset}(L^S(Y))$, its values when F is a *pair set* with $\#F = 2$ determine a strict preference relation \succ_C , a strict dispreference relation \prec_C , and an indifference relation \sim_C . These three binary relations are defined so that for each pair $\lambda^S, \mu^S \in L^S(Y)$, one has

$$\lambda^S \left\{ \begin{array}{l} \succ_C \\ \sim_C \\ \prec_C \end{array} \right\} \mu^S \quad \text{according as} \quad C(\{\lambda^S, \mu^S\}) = \left\{ \begin{array}{l} \{\lambda^S\} \\ \{\lambda^S, \mu^S\} \\ \{\mu^S\} \end{array} \right\} \quad (41)$$

Underlying the choice function $F \mapsto C(F)$ specified by (40), there is a single corresponding binary weak preference relation \succeq_C on $L^S(Y)$, called the *base relation*. For each pair $\lambda^S, \mu^S \in L^S(Y)$, this base relation satisfies

$$\lambda^S \succeq_C \mu^S \iff \lambda^S \in C(\{\lambda^S, \mu^S\}) \iff \lambda^S \succ_C \mu^S \text{ or } \lambda^S \sim_C \mu^S \quad (42)$$

Finally, we mention the corresponding weak dispreference relation \preceq_C defined on $L^S(Y)$ so that

$$\lambda^S \preceq_C \mu^S \iff \mu^S \in C(\{\lambda^S, \mu^S\}) \iff \lambda^S \prec_C \mu^S \text{ or } \lambda^S \sim_C \mu^S \quad (43)$$

Evidently both the weak preference relation \succeq_C and the weak dispreference relation \preceq_C are *complete* in the sense that, for each pair $\lambda^S, \mu^S \in L^S(Y)$, one has:

1. either $\lambda^S \succeq_C \mu^S$ or $\mu^S \succeq_C \lambda^S$ or both;
2. either $\lambda^S \preceq_C \mu^S$ or $\mu^S \preceq_C \lambda^S$ or both.

6.4 Bayesian Rationality and Expected Utility

Let $F \mapsto C(F)$ be any choice function satisfying (40) and $C(F) \subseteq F$ for all $F \in \mathcal{F}_{\setminus \emptyset}(L^S(Y))$ that corresponds to the base preference relation \succsim_C defined on the space $L^S(Y)$ of AA lotteries λ^S by (41). Then the mapping $L^S(Y) \ni \lambda^S \mapsto U^S(\lambda^S) \in \mathbb{R}$ is a *utility function* which *represents* the preference relation \succsim_C on $L^S(Y)$ just in case, for all $\lambda^S, \mu^S \in L^S(Y)$, one has

$$\lambda^S \succsim_C \mu^S \iff U^S(\lambda^S) \geq U^S(\mu^S) \quad (44)$$

Given the specified non-empty finite set S of uncertain states of the world, an *interior subjective probability mass function* is a mapping $S \ni s \mapsto \mathbb{P}(s) \in (0, 1]$ that satisfies $\sum_{s \in S} \mathbb{P}(s) = 1$. As discussed in Section 5.4 of Hammond (1998b) and in Hammond (2022), the restriction to positive probabilities is to avoid the difficulties that arise in continuation subtrees of a decision tree when zero probabilities are allowed.

Then the choice function $F \mapsto C(F)$ on the domain $\mathcal{F}_{\setminus \emptyset}(L^S(Y))$, together with the associated base preference relation \succsim_C on the space $L^S(Y)$, are both *Bayesian rational* just in case there exist a probability mass function $S \ni s \mapsto \mathbb{P}(s) \in (0, 1]$, as well as a *Bernoulli utility function* $Y \ni y \mapsto u(y) \rightarrow \mathbb{R}$, such that \succsim_C is represented, in the sense that (44) is satisfied, by the *von Neumann subjective expected utility function* defined for all AA lotteries $\lambda^S = \langle \lambda_s \rangle_{s \in S} \in L^S(Y)$ by the double sum

$$U^S(\lambda^S) = \sum_{s \in S} \mathbb{P}(s) \sum_{y \in Y} \lambda_s(y) u(y) \quad (45)$$

6.5 Normalized Utility

Recall that, as explained in Section 6.1, for each $y \in Y$, we use δ_y to denote the unique *degenerate* probability measure in $\Delta(Y)$ that satisfies $\delta_y(\{y\}) = 1$. To avoid trivialities, we assume that there exist at least three consequences $\underline{y}, y^0, \bar{y}$ in the domain Y such that, for the three corresponding degenerate lotteries $\delta_{\bar{y}}, \delta_{y^0}, \delta_{\underline{y}}$, the base strict preference relation \succ satisfies the strict preference property $\delta_{\bar{y}} \succ \delta_{y^0} \succ \delta_{\underline{y}}$.

Consider any Bernoulli utility function $Y \ni y \mapsto u(y) \rightarrow \mathbb{R}$ and the associated preference relation \succsim on $\Delta(Y)$ that satisfies

$$\lambda \succsim \mu \iff U(\lambda^S) \geq U(\mu^S) \quad (46)$$

for the *von Neumann objective expected utility function* on $\Delta(Y)$ defined by

$$U(\lambda) = \sum_{y \in Y} \lambda(y)u(y) \quad (47)$$

As explained in Hammond (1998a, Section 2.3), assuming Bayesian rationality, the ratio $\frac{u(y^0) - u(\underline{y})}{u(\bar{y}) - u(\underline{y})}$ of utility differences equals, in economists' terminology, the constant marginal rate of substitution along an indifference curve between shifts in probability: (i) from consequence \underline{y} to y^0 ; (ii) from consequence \underline{y} to \bar{y} . This leads us to say that two Bernoulli utility functions $y \mapsto u(y)$ and $y \mapsto \tilde{u}(y)$ are *equivalent* just in case, for every triple $\underline{y}, y^0, \bar{y}$ of consequences in Y satisfying $\delta_{\bar{y}} \succ \delta_{y^0} \succ \delta_{\underline{y}}$ and so $u(\bar{y}) > u(y^0) > u(\underline{y})$, the corresponding ratios of utility differences satisfy

$$\frac{u(y^0) - u(\underline{y})}{u(\bar{y}) - u(\underline{y})} = \frac{\tilde{u}(y^0) - \tilde{u}(\underline{y})}{\tilde{u}(\bar{y}) - \tilde{u}(\underline{y})} \quad (48)$$

But (48) holds for every triple $\underline{y}, y^0, \bar{y}$ satisfying $u(\bar{y}) > u(y^0) > u(\underline{y})$ if and only if there exist an additive constant $\alpha \in \mathbb{R}$ and a positive multiplicative constant $\rho \in \mathbb{R}$ such that, for all $y \in Y$, one has

$$\tilde{u}(y) = \alpha + \rho u(y) \quad (49)$$

Now, given any pair \underline{u}, \bar{u} of real numbers with $\bar{u} > \underline{u}$ and any Bernoulli utility function $Y \ni y \mapsto u(y) \rightarrow \mathbb{R}$, there exist two unique constants α and $\rho > 0$ such that the transformed utility function defined by (49) is a *normalized* utility function that satisfies $\tilde{u}(\underline{y}) = \underline{u}$ and $\tilde{u}(\bar{y}) = \bar{u}$. Indeed the two constants we need are given by

$$\rho = \frac{\bar{u} - \underline{u}}{u(\bar{y}) - u(\underline{y})} \quad \text{and then} \quad \alpha = \bar{u} - \rho u(\bar{y}) = \underline{u} - \rho u(\underline{y}) \quad (50)$$

From now on let u denote the unique Bernoulli utility function $Y \ni y \mapsto u(y) \rightarrow \mathbb{R}$ whose expected values defined by (47) satisfy (46), and which has been normalized to satisfy

$$u(\underline{y}) = \underline{u} \quad \text{and} \quad u(\bar{y}) = \bar{u} \quad (51)$$

6.6 Finite Decision Trees and Their Continuations

In mathematical terminology a *graph* (N, E) is a non-empty set N of vertices or *nodes* n , together with a specified set $E \subseteq N \times N$ of *edges* $e = (n, n')$ which connect a subset of disjoint pairs $(n, n') \in N \times N$ of nodes. The graph (N, E) is *finite* just in case the set N of nodes is finite, which implies that the set E of edges is also finite.

The graph (N, E) is *directed* just in case there is an antisymmetric and complete binary relation $>_{+1}$ on N with the property that for each edge $(n, n') \in E$, either $n >_{+1} n'$, or $n' >_{+1} n$, but not both. In the case of a directed tree, we can write each edge (n, n') of E in the form $n \rightarrow n'$, where $n' >_{+1} n$. The sequence $(n_1, n_2, \dots, n_\ell)$ of ℓ nodes in N is a *path* of length $\ell \in \mathbb{N}$ in the directed graph (N, E) just in case $n_{k+1} >_{+1} n_k$ for $k = 1, 2, \dots, \ell - 1$.

The graph (N, E) is a *rooted directed tree* just in case there is a unique *initial node* $n_0 \in N$ (or root, or seed, or entry point) such that, for every other node $n \in N \setminus \{n_0\}$ of the graph, there is a unique path $(n_0, n_1, n_2, \dots, n)$ which starts at node n_0 and ends at node n .

Definition 1. Let $T = (N, E)$ denote any finite rooted directed tree.

1. Given any pair of nodes $n, n' \in N$, say that $n' \in N$ (ultimately) succeeds n just in case there is a path $(n_1, n_2, \dots, n_\ell)$ of nodes in N of length ℓ which joins $n_1 = n$ to $n_\ell = n'$. Let $N_{>n}$ denote the set of all nodes $n' \in N$ that succeed n .
2. For each $n \in N$, define the continuation subtree $T_{\geq n} = (N_{\geq n}, E_{\geq n})$ whose initial node is n as the unique tree in which:
 - $N_{\geq n} = \{n\} \cup N_{>n}$;
 - $E_{\geq n}$ is the restriction to $E \cap (N_{\geq n} \times N_{\geq n})$ of edges in E .
3. Say that any other node $n' \in N$ immediately succeeds n just in case the ordered pair (n, n') or $n \rightarrow n'$ is a directed edge of T .
4. Denote the set of all the immediate successors of n by

$$N_{\geq n}^{+1} := \{n' \in N \mid (n, n') \in E\} = \{n' \in N \mid n' >_{+1} n\} \quad (52)$$

Recall from Section 1.1 the distinction between roulette and horse lotteries due to Anscombe and Aumann (1963). The following definition builds on those in Hammond (1988b; 1998a, b; 2022).

Definition 2. Given the non-empty consequence domain Y and non-empty finite set S of possible states of the world, define $\mathcal{T}^S(Y)$ as the collection of all finite decision trees which combine a rooted directed tree $T = (N, E)$ with an event correspondence $N \ni n \mapsto S_{\geq n} \in 2^S \setminus \{\emptyset\}$ specifying what set $S_{\geq n}$ of states are possible in each continuation subtree $T_{\geq n}$. Moreover, the set N of nodes in the decision tree is partitioned into four pairwise disjoint subsets:

1. the set N^d of decision nodes n at which: (i) at n the agent must choose, for some node n' in the set $N_{\geq n}^{+1}$ of immediate successors defined in (52), the directed edge $n \rightarrow n'$; (ii) $S_{\geq n'} = S_{\geq n}$ for all $n' \in N_{\geq n}^{+1}$;
2. the set N^c of chance nodes n where: (i) a directed edge $n \rightarrow n'$ emanating from n is determined randomly by a specified roulette lottery $N_{\geq n}^{+1} \ni n' \mapsto \pi(n'|n) \in (0, 1]$; ¹⁷ (ii) $S_{\geq n'} = S_{\geq n}$ for all $n' \in N_{\geq n}^{+1}$;
3. the set N^e of event nodes n where a directed edge $n \rightarrow n'$ emanating from n is determined by a horse lottery whose outcome partitions the event $S_{\geq n}$ into the collection $\{S_{\geq n'} \mid n' \in N_{\geq n}^{+1}\}$ of non-empty pairwise disjoint sub-events $S_{\geq n'}$ which, for each $n' \in N_{\geq n}^{+1}$, consist of states $s \in S$ that can occur in the subtree $T_{\geq n'}$;
4. the non-empty set N^t of terminal nodes n at which $N_{\geq n}^{+1} = \emptyset$, so no edge emanates, and which are each mapped to an Anscombe–Aumann consequence lottery $\gamma(n) = \langle \gamma_s \rangle_{s \in S_{\geq n}} \in L^{S_{\geq n}}(Y)$ whose outcomes y belong to the specified consequence domain Y .

Say that a decision tree $T \in \mathcal{T}^S(Y)$ is:

- *deterministic* just in case $N^c = N^e = \emptyset$;
- *risky* just in case $N^c \neq \emptyset$ but $N^e = \emptyset$; ¹⁸
- a *Savage tree* just in case $N^e \neq \emptyset$ but $N^c = \emptyset$;
- an *Anscombe–Aumann tree* just in case $N^c \neq \emptyset$ and $N^e \neq \emptyset$.

¹⁷See Hammond (1988b) for an explanation of why, if there is a chance node $n \in N^c$ and a node $n' \in N_{\geq n}^{+1}$ at which $\pi(n'|n) = 0$, then all consequence lotteries must be indifferent.

¹⁸Raiffa (1968) focused on risky decision trees with *pecuniary consequences* in the form of payoffs measured in dollars.

6.7 Evaluations of Continuation Subtrees

Consider the orthodox “unenlivened” decision model which is represented by any finite decision tree T in the domain $\mathcal{T}^S(Y)$ of trees with non-empty finite state space S and AA lottery consequences in the domain $\cup_{S' \in 2^S \setminus \{\emptyset\}} L^{S'}(Y)$. Working backwards within each decision tree, as usual in dynamic programming, it is possible in principle to use a recursive procedure in order to calculate the *evaluation* $v(T_{\geq n})$ as the subjectively expected continuation value of reaching any node $n \in N$, which is the initial node of the continuation subtree $T_{\geq n}$. The details of this backward recursive procedure are described in the remainder of this section.

In any finite decision tree T , the backward recursion starts at any terminal node $n \in N^t$. As discussed in Section 6.6, the specified consequence of reaching any terminal node $n \in N^t$ of tree T is the AA consequence lottery $\gamma(n) = \langle \gamma_s \rangle_{s \in S_{\geq n}} \in L^{S_{\geq n}}(Y)$. Then the evaluation $v(T_{\geq n})$ of the subtree $T_{\geq n}$, whose only node is the terminal node n , is the expected utility of the AA consequence lottery $\gamma(n)$, which is specified by (45) as the double sum

$$v(T_{\geq n}) = U^{S_{\geq n}}(\gamma(n)) = \sum_{s \in S_{\geq n}} \mathbb{P}(s) \sum_{y \in Y} \gamma_s(y) u(y) \quad (53)$$

At any non-terminal node $n \in N \setminus N^t$, the evaluation $v(T_{\geq n})$ of the continuation subtree $T_{\geq n}$ starting at the initial node n depends upon the set $\{v(T_{\geq n'}) \mid n' \in N_{\geq n}^{+1}\}$ of evaluations of all the continuation subtrees starting at a node $n' \in N_{\geq n}^{+1}$ which immediately succeeds n . Specifying the correct formula for $v(T_{\geq n})$ requires considering separately three cases, depending upon whether n is a chance, event, or decision node.

In the first case when n is a chance node whose immediate successors $n' \in N_{\geq n}^{+1}$ occur with respective specified positive probabilities $\pi(n'|n)$, the relevant recursion takes the obvious form

$$v(T_{\geq n}) = \sum_{n' \in N_{\geq n}^{+1}} \pi(n'|n) v(T_{\geq n'}) \quad (54)$$

The second case occurs when n is an event node, each of whose immediate successors $n' \in N_{\geq n}^{+1}$ determines which is the relevant cell of the partition $\{S_{\geq n'} \mid n' \in N_{\geq n}^{+1}\}$ of the event $S_{\geq n}$ into pairwise disjoint sets. In this case the objectively specified probabilities $\pi(n'|n)$ that appear in (54) need to be replaced by subjective conditional probabilities $p(n'|n)$ derived from the relevant subjective probabilities $\mathbb{P}(s)$ for different states $s \in S_{\geq n}$. Because of

our requirement that $\mathbb{P}(s) > 0$ for all $s \in S$, these conditional probabilities $p(n'|n)$ are all well defined, and can be calculated as

$$p(n'|n) = \sum_{s \in S_{\geq n'}} \mathbb{P}(s) / \sum_{s \in S_{\geq n}} \mathbb{P}(s) \quad (55)$$

So, when n is an event node with subjective probabilities $p(n'|n)$ given by (55) rather than a chance node with hypothetical or objective probabilities $\pi(n'|n)$, the previous formula (54) is changed to

$$v(T_{\geq n}) = \sum_{n' \in N_{\geq n}^{+1}} p(n'|n) v(T_{\geq n'}) \quad (56)$$

In the third and final case when n is a decision node, we apply the standard *optimality principle* of stochastic dynamic programming. This requires any current *optimal decision* $n^* \in N_{\geq n}^{+1}$ to be the first step toward achieving the highest possible expected value resulting from an appropriate plan for all subsequent decisions. Consider the induction hypothesis that, for each node $n' \in N_{\geq n'}^{+1}$, the value $v(T_{\geq n'})$ is the maximum possible evaluation the agent can achieve by choosing an optimal decision at each decision node of $T_{\geq n'}$. This is trivially true when n' is a terminal node, so there is no decision to make at node n' . If this hypothesis is true at each node $n' \in N_{\geq n}^{+1}$, then any optimal decision at node n must be to move along an edge $n \rightarrow n^*$ to an immediately succeeding node n^* which maximizes the evaluation $v(T_{\geq n'})$ with respect to n' subject to $n' \in N_{\geq n}^{+1}$, where the set $N_{\geq n}^{+1}$ is finite. In other words, one must satisfy

$$n^* \in \arg \max_{n' \in N_{\geq n}^{+1}} v(T_{\geq n'}) \quad (57)$$

So the appropriate recursion when n is a decision node is

$$v(T_{\geq n}) = v(T_{\geq n^*}) = \max_{n' \in N_{\geq n}^{+1}} v(T_{\geq n'}) \quad (58)$$

Together, therefore, the four equations (53), (54), (56), and (58) do indeed determine $v(T_{\geq n})$ by backward recurrence in the four different cases.

7 Truncated Decision Trees

7.1 Bounded Rationality as Bounded Modelling

The conditionally expected enlivened evaluation associated with entering any continuation subtree $T_{\geq n}$ of a decision tree $T \in \mathcal{T}^S(Y)$ is $v(T_{\geq n})$. In principle, this can be calculated by following the procedure of backward recursion

that was set out in Section 6.7. Carrying out all the required steps of the computation by backward recursion, however, obviously becomes much more challenging, if not practically impossible, as at least some of the continuation subtrees $T_{\geq n}$ become more complicated.

As discussed in Section 1.2, these practical limitations on being able to calculate an optimal decision may force the decision maker to resort to some kind of bounded model. In this section, we consider bounded models that result when the original decision tree gets truncated by having some entire continuation subtrees removed. In the game of Chess, for example, which was considered in Section 5, this is the kind of tree that results when a player looks at most k moves ahead, for some small integer $k \in \mathbb{N}$.

7.2 Cuts and Truncation Nodes

The following definition captures the idea that any truncation of a decision tree will occur at a truncation node or cut. This cut node has the property that all strictly succeeding nodes are removed, so the truncation node becomes a terminal node.

Definition 3. *Let T be any decision tree in the domain $\mathcal{T}^S(Y)$. The tree $\hat{T} = (\hat{N}, \hat{E})$ is a truncation of T just in case there is a set X of one or more cuts or non-terminal truncation nodes $x \in N \setminus N^t$ such that:*

1. *there is a pairwise disjoint collection $\{N_{>x} \mid x \in X\}$ of removed sets $N_{>x}$ of non-initial nodes $n \in N \setminus \{n_0\}$, each associated with a continuation subtree $T_{\geq x}$ whose initial node is at a cut $x \in X$;*
2. *the truncated set \hat{N} of remaining nodes is $N \setminus \cup_{x \in X} N_{>x}$;*
3. *the truncated set \hat{E} of remaining directed edges is the restriction $(\hat{N} \times \hat{N}) \cap E$ to $\hat{N} \times \hat{N}$ of the original set $E \subset N \times N$ of directed edges $n \rightarrow n'$ in T ;*
4. *each cut $x \in X$ is a terminal node $x \in \hat{N}^t$ of the truncated tree \hat{T} .*

7.3 Subjective Evaluations in Truncated Trees

Given a finite decision tree $T = (N, E)$, in order to make the truncation $\hat{T} = (\hat{N}, \hat{E})$ of T a decision tree, each terminal node $n \in \hat{N}^t$ must be assigned

a consequence $\hat{\gamma}(n)$. In case $n \in \hat{N}^t \cap N^t$, which is when the terminal node n of \hat{T} happens also to be a terminal node of T , the consequence $\hat{\gamma}(n)$ of n in \hat{T} should be the same as its consequence in T , implying that $\hat{\gamma}(n) = \gamma(n)$.

In case $n \in \hat{N}^t \setminus N^t$, however, no consequence $\hat{\gamma}(n)$ has yet been specified at node n , which has become a terminal node of \hat{T} but not of T . In this case we assign to the cut node n a new *subjective evaluation* $\hat{\gamma}(n)$. This is not a consequence, or even a consequence lottery. Instead it is assumed to be a real number in the convex hull of the range $u(Y)$ of the normalized utility function $Y \ni y \mapsto u(y) \in \mathbb{R}$ that was constructed in Section 6.5. This construction implies that each possible subjective evaluation $\hat{\gamma}(n)$ is the expected utility of some lottery in the domain $\Delta(Y)$ of finitely supported lotteries over the consequence domain Y .

In the full tree T , with unbounded rationality, an appropriate evaluation $v(T_{\geq n})$ of the node $n \in \hat{N}^t \setminus N^t$ should be an expected utility level that is calculated by backward recursion, as in Section 6.7. In the truncated tree \hat{T} , however, the best we can do is to use an estimate $\hat{\gamma}(n) \in \mathbb{R}$ of what expected utility $v(T_{\geq n})$ would have resulted had the agent been able to calculate this in the full continuation subtree $T_{\geq n}$, without any truncation.

In the rest of the paper, to allow for truncated trees, we distinguish between:

1. “consequence” terminal nodes n , which each have a specified AA lottery consequence $\gamma(n) \in L^{S_{\geq n}}(Y)$ attached;
2. “evaluation” terminal nodes n , which each have a subjective evaluation $\hat{\gamma}(n) \in \mathbb{R}$ attached.

7.4 Cardinally Equivalent Evaluations

Suppose that the strictly increasing affine transformation given by $\tilde{u}(y) = \alpha + \rho u(y)$ for all $y \in Y$, with $\rho > 0$, is applied to the utility function $Y \ni y \mapsto u(y) \rightarrow \mathbb{R}$ that appears in (53), resulting in the cardinally equivalent function $y \mapsto \tilde{u}(y)$. This transformation can be achieved by replacing the lower constant \underline{u} and upper constant \bar{u} in the normalization condition (51) by their respective transformed values $\alpha + \rho \underline{u}$ and $\alpha + \rho \bar{u}$.

In decision trees that are not truncated, it is obvious and easy to show that if the backward recursion procedure described in Section 6.7 is applied to the transformed utility function, then all the subtree evaluation functions

$N \ni n \mapsto v(T_{\geq n}) \in \mathbb{R}$ undergo the equivalent strictly increasing affine transformation given by $\tilde{v}(T_{\geq n}) = \alpha + \rho v(T_{\geq n})$ for all $n \in N$, so that subtree evaluations are also cardinally equivalent. It is also obvious that this common affine transformation of the utility and subtree evaluation functions has no effect on the optimal decision at each decision node which is given by (57).

In order to preserve these properties in a truncated decision tree, it is clear that the subjective evaluation $\hat{\gamma}(n)$ at any node $n \in \hat{N}^t \setminus N^t$ where a truncation occurs should also undergo the same strictly increasing affine transformation. That is, we should replace each subjective evaluation $\hat{\gamma}(n)$ by the transformed subjective evaluation $\tilde{\gamma}(n) = \alpha + \rho \hat{\gamma}(n)$.

Recall from Section 7.3 that $\hat{\gamma}(n)$ should be an estimate of the expected utility from following an optimal policy in the continuation decision tree $T_{\geq n}$ which, apart from its initial node n , is removed by truncation. So when the utility function undergoes a strictly increasing affine transformation, these transformed subjective evaluations $\tilde{\gamma}(n) = \alpha + \rho \hat{\gamma}(n)$ are the obvious revisions of these estimated expected utilities.

7.5 Monte Carlo Tree Search

Long before computer games became popular recreations, mathematicians viewed games as models of decision making. The general understanding of decisions, however, has been impeded by the ambiguity of some of the basic components of game-tree search. In particular, the static evaluation function, or determination of a node's merit based on directly detectable features, has never been adequately defined. The expected-outcome model proposes that the appropriate value to assign a node is the expected value of a game's outcome given random play from that node on.

— from the abstract to Bruce Abramson's (1987) Ph.D. dissertation, eventually published as Abramson (1991).

We turn next to a powerful kind of computer algorithm that has proved extremely fruitful in estimating continuation values for at least some particular kinds of truncated game. Indeed, exploiting Abramson's key idea of Monte Carlo tree search (MCTS) described in the quotation above helped to inspire a generation of computer programs that:

1. in the case of Chess, led to the Stockfish software engine that would easily beat any human player over any sufficiently long run of games;
2. but in the case of Go, was unable to defeat the best human players.

Eventually, this generation of algorithms became superseded by AlphaZero, based on the kind of artificial neural network that has become a key part of what has come to be known as “artificial intelligence”. Such algorithms have proved far better at playing Chess than programs of the Stockfish generation, while finally becoming able to beat the best human players at Go. The key paper by Silver et al. (2018), however, reports that MCTS remains part of AlphaZero.

In decision trees, MCTS evidently requires evaluating any continuation decision tree by calculating the normalized expected utility from a “Monte Carlo” sample of simulations in which the agent makes a suitably randomized choice at each decision node in that continuation.

8 Enlivened Decision Trees

8.1 Simple Enlivenment as Enrichment

It is time to return to the main task of the paper. This is to explain how a relevant version of the Bayesian rationality hypothesis can be applied even to enlivened decision trees. We do this not only for entire decision trees T , but also for any continuation decision tree $T_{\geq n}$, as defined in Section 6.6. Since each node n of a tree T is the initial node of $T_{\geq n}$, this helps us to focus, without loss of generality, on the initial node of any decision tree we consider.

Given any decision tree $T \in \mathcal{T}^S(Y)$, we begin by defining a “simple enlivenment” as a suitable modification of any particular “unenlivened” continuation subtree $T_{\geq n}$ whose initial node is n . This minimal enlivenment will then result from appending to $T_{\geq n}$ an appropriate finite set of extra nodes which are initial nodes of corresponding continuation subtrees. This, after all, is what happened when we described the enlivened decision trees in the three examples of Sections 3, 4, and 5.

Indeed, in the model of Odysseus and the Sirens set out in Section 3, there were three stages of enlivenment. These came from adding new nodes one step at a time through the succession of four trees illustrated in Figures

1–4. Second, in Section 4, the enlivened two-period linear-quadratic portfolio problem which is described in Section 4.5 is the result of adding extra dimensions to the original problem that was set out in Section 4.1. Third, the Chess “blunder of the century” described in Section 5.3 seems to have resulted from Kramnik as Black calculating the likely consequences of his planned moves using a bounded model which excludes the move 35. Qe4–h7 that Deep Fritz, playing White, eventually used to deliver checkmate. In this sense, Deep Fritz was using an enlivened version of the bounded model which Kramnik had been using, with 35. Qe4–h7 checkmate as an extra possible move which is included in the enriched model, but not in Kramnik’s bounded model.

8.2 Limits to Enlivenment

Suppose that the agent’s original unenlivened continuation subtree $T_{\geq n}$ is minimally enlivened to a new enlivened subtree $T_{\geq n}^+$ which extends $T_{\geq n}$ by recognizing previously excluded possibilities. Consider any new information which, when the agent is at a decision node $n \in N^d$, arrives in time to change the modelled feasible set $M_n := N_{\geq n}^{+1}$ of possible immediate successors of n . These are the nodes n_m which, for some $m \in M_n$, can be reached by moving along the edge $n \rightarrow n_m$ that emanates from the initial node n of $T_{\geq n}$. As a result, the set M_n is replaced by a new expanded feasible set M_n^+ which is incorporated in a new tree T^+ . We argue that this expansion is really a trivial enlivenment of T because, before making a decision at node n , the agent has time to recognize that the feasible set of moves $n \rightarrow n_m$ at node n is M_n^+ rather than only the proper subset M_n . And this, of course, is exactly what the agent should do.

Thus, a necessary condition for the agent to be able to decide to move to a node n_m with $m \in M_n^+$ in the expanded continuation subtree $T_{\geq n}^+$ is that node n_m must be included in the agent’s original continuation subtree $T_{\geq n}$. This is also sufficient; if the node n_m and edge $n \rightarrow n_m$ are in the tree $T_{\geq n}^+$, then the agent, while still at node n in the tree, could in principle choose the move $n \rightarrow n_m$ to the node n_m . Accordingly, we assume that M_n remains unaffected by any enlivenment of the continuation subtree $T_{\geq n}$. Only after a chosen decision node $n_m \in M_n$ has already been reached can any enlivening of the initial continuation subtree $T_{\geq n}$ occur.

Similarly, even if the initial node n of the original continuation subtree $T_{\geq n}$ is not a decision node, we still assume that simple enlivenment can never

occur *at* node n . Instead, it must occur at a new node which belongs to $T_{\geq n}^+$ but not to $T_{\geq n}$.¹⁹

8.3 Influence Diagrams as Compressed Decision Trees

We have assumed that any decision tree $T \in \mathcal{T}^S(Y)$ is finite, as of course is the continuation subtree $T_{\geq n}$ for any $n \in N$. Nevertheless, each continuation subtree can be arbitrarily large, as can the set M_n of moves $n \rightarrow n'$ in E that are possible at each node n , whether node n is a decision, chance, or event node. So generally it is impossible to represent all the nodes in a continuation subtree $T_{\geq n}$ in a single diagram of that subtree.

Accordingly, instead of trying to represent graphically any full continuation subtree $T_{\geq n}$, we adapt the standard method described by Howard and Matheson (2005). This uses “influence diagrams” in order to compress what would otherwise be the often extraordinarily convoluted full graphical representation of a decision tree. In the simple cases economists usually consider, they often describe influence diagrams as showing “time lines”, each of which represents the typical member of what may be a large set of several paths through the tree. Instead of ending in one specific consequence, such influence diagrams will usually end, for each terminal node $n \in N^t$, in a variable consequence $\gamma(n) \in L^{S_{\geq n}}(Y)$.




Figure 6: Reduced Influence Diagram for the Move m

The adaptation we consider reduces the decision tree even further to that illustrated in Figure 6. Whereas an influence diagram typically shows an entire typical path in a tree, right through to a terminal node with its consequence, Figure 6 shows only one typical edge $n \rightarrow n_m$. This corresponds to a typical move to n_m where m belongs to the set M_n of moves which are possible at the initial node n of the continuation subtree $T_{\geq n}$. Figure 6 also shows the node n_m with a label indicating the new continuation subtree $T_{\geq n_m}$ whose initial node n_m is reached after the move m .

¹⁹A referee has suggested allowing enlivenment to occur at existing nodes in a decision tree. This, however, would require specifying the effect of enlivenment for the three separate cases of decision, chance, and event nodes. It seems simpler to introduce an extra enlivenment node immediately prior to any existing node at which enlivenment is thought to occur.

8.4 Tree Stems as Enlivenment Edges

For simplicity, the initial focus is on a minimal but prominent kind of simple enlivenment. Given a decision tree $T = (N, E)$ and a particular node n , consider the corresponding continuation subtree $T_{\geq n}$ of the tree T , with initial node n and finite graph $(N_{\geq n}, E_{\geq n})$. As in Section 8.1, let M_n denote the set of all possible moves m along a directed edge $n \rightarrow n_m$ in $E_{\geq n}$, starting at the initial node of $T_{\geq n}$. Then the minimal enlivenment we consider starts by affecting only one particular directed edge $n \rightarrow n_m$ of $E_{\geq n}$. The enlivenment starts at a *bud* or *pre-enlivenment node* e_m^- which, as shown in Figure 7, is inserted into the middle of the corresponding directed edge $n \rightarrow n_m$ of the original subtree $T_{\geq n}$. The first edge of the enlivenment is the *stem* or *enlivenment edge* (e_m^-, e_m^+) that joins the bud e_m^- to the later *post-enlivenment node* e_m^+ , which is the initial node of the *enlivening subtree* $T_{\geq e_m^+}^+$.




Figure 7: A Bud and Stem with an Enlivened Continuation

To complete the specification of the minimally enlivened tree T^+ , we make the bud e_m^- an event node where the outcome of a horse lottery determines whether an uncertain deviation from the path $n \rightarrow e_m^- \rightarrow n_m$ back to the original tree T does or does not occur. Specifically:

1. First, if no deviation occurs, then the relevant immediate successor of e_m^- in the minimal enlivenment $T_{\geq n}^+$ of $T_{\geq n}$ is the initial node n_m of a copy of the original unenlivened continuation subtree $T_{\geq n_m}$ that emanates from n_m in the original unenlivened tree T .
2. Second, if a deviation does occur, then the relevant immediate successor of e_m^- is the *post-enlivenment node* e_m^+ . This is the initial node of a finite unenlivened continuation subtree $T_{\geq e_m^+}$ that gets appended to $T_{\geq n}$ at e_m^+ in the process of enlivening $T_{\geq n}$ to the new tree $T_{\geq n}^+$.

Now we move on to consider simply rather than minimally enlivened decision trees. By definition, the new tree $T^+ = (N^+, E^+)$ is a *simple enlivenment* of the original decision tree $T = (N, E)$ just in case it is the result of a finite

sequence of minimal enlivenments of T . For each node $n \in N$, recall the notation $M_n = N_{\geq n}^{+1}$ for the set of nodes that immediately succeed n . Then $T^+ = (N^+, E^+)$ is a simple enlivenment of $T = (N, E)$ just in case, for each $n \in N$, there is a non-empty subset M_n^+ of M_n with the property that, for each $m \in M_n^+$, there is a bud e_m^- inserted into the edge $n \rightarrow n_m$ along with a stem in the form of an extra enlivenment edge $e_m^- \rightarrow e_m^+$ ending in the initial node of the new continuation subtree $T_{\geq e_m^+}^+$.

8.5 Enlivened Consequences and States

Eventually we will explore the implications of Bayesian rational behaviour for the entire family of simply enlivened decision trees T^+ . Before doing so, however, we need to recognize that, even in the relatively simple examples of Sections 3 and 4, enlivenment may expand: (i) the consequence domain Y of the Bernoulli utility function $Y \ni y \mapsto u(y) \rightarrow \mathbb{R}$, which has been normalized to satisfy (51); (ii) the finite space S of uncertain states of the world.

Indeed, for each move $m \in M_n^+$ at which the possibility of enlivenment is modelled, and so for each corresponding continuation subtree $T_{\geq e_m^+}^+$, there is a finite enriched domain Y_m^+ of consequences y and an enriched domain S_m^+ of possible states s that could feasibly result from following an appropriate path from e_m^+ to a terminal node. Then the *enlivened consequence domain* and *enlivened state space* are defined as the unions

$$Y^+ = \cup_{m \in M_n^+} Y_m^+ \quad \text{and} \quad S^+ = \cup_{m \in M_n^+} S_m^+ \quad (59)$$

over the finite collection of, respectively, all possible enriched consequence domains Y_m^+ and all possible enriched state spaces S_m^+ . Let $\mathcal{T}^{S^+}(Y^+)$ denote the relevant domain of possible enlivened decision trees with the enlived consequence domain and state space. Let $L^{S^+}(Y^+)$ denote the relevant space of possible enlivened AA consequence lotteries.

Recall that the two consequences $\underline{y}, \bar{y} \in Y$ were used in Section 6.5 in order to construct the unique normalized utility function that satisfies (51). Notice that the new domain $\mathcal{T}^{S^+}(Y^+)$ of decision trees must include, for each consequence $y \in Y^+$, the set $\mathcal{T}(\Delta(\{y, \bar{y}, y\}))$ of risky finite decision trees, without any event nodes, horse lotteries, or uncertain states, whose roulette consequence lotteries are restricted to probability mixtures of the three degenerate lotteries $\delta_{\underline{y}}$, $\delta_{\bar{y}}$, and δ_y . Then the construction set out in Section 6.5 can be repeated to determine the utility $u(y^+)$ of each extra

consequence $y^+ \in Y^+ \setminus Y$. The result is a normalized enlivened Bernoulli utility function

$$Y^+ \ni y^+ \mapsto u^+(y^+) \rightarrow \mathbb{R} \quad (60)$$

that satisfies (51) not only for consequences y in the original domain Y , but also for any extra consequences y^+ in the entire extended domain Y^+ . Moreover, the expectation of u^+ defined by (60) will represent the agent's extended preference relation \succsim^+ on the whole lottery domain $\Delta(Y^+)$. Also, because the two particular consequences \underline{y}, \bar{y} used in the earlier normalization (51) are in Y^+ as well as in Y , we can impose the obvious counterpart

$$u^+(\underline{y}) = \underline{u} \quad \text{and} \quad u^+(\bar{y}) = \bar{u} \quad (61)$$

of that earlier normalization. The resulting function (60) will then extend $Y \ni y \mapsto u(y) \rightarrow \mathbb{R}$ to the expanded domain Y^+ .

Similarly, by considering the domain $\mathcal{T}^{S^+}(Y^+)$ instead of $\mathcal{T}^S(Y^+)$, the construction of the subjective probabilities $\mathbb{P}(s)$ of each state $s \in S$ can be extended to a construction of the enlivened subjective probabilities $\mathbb{P}^+(s)$ of each state $s \in S^+$.

8.6 Bayesian Rationality in Minimally Enlivened Trees

Let $T = (N, E)$ denote any decision tree in the domain $\mathcal{T}^S(Y)$. Consider once again a minimal enlivenment, as set out in Section 8.4, with the property that there exists a unique edge $n \rightarrow n_m$ in E into which a unique bud e_m^- is inserted. This insertion was illustrated in Figure 7. Note that the continuation subtrees $T_{\geq e_m^-}^+$ and $T_{\geq e_m^+}^+$ both belong to the enlivened domain $\mathcal{T}^{S^+}(Y^+)$ of finite decision trees.

Consider next the subtree evaluations that were constructed in Section 6.7 by backward recurrence, using whichever of the four equations (53), (54), (56), and (58) applies to each successive node. Bayesian rationality implies that exactly the same construction can be applied in order to calculate the subjective evaluation $v^+(T_{\geq e_m^+}^+)$ of the enlivened continuation subtree $T_{\geq e_m^+}^+$.

To continue the backward recursion, we also need to go back one more step in order to construct an appropriate evaluation $v^+(T_{\geq e_m^-}^+)$ at the bud e_m^- of the enlivened continuation decision tree $T_{\geq e_m^-}^+$. In order to do so, we apply the refined form of Bayesian rationality specified in Hammond (2022). It implies that there exists a unique positive subjective probability $\eta_m \in (0, 1)$

that enlivenment occurs at the event node e_m^- , and so a positive probability $1 - \eta_m \in (0, 1)$ that enlivenment does not occur at e_m^- . The argument for insisting that $0 < \eta_m < 1$ is that if $\eta_m = 0$, then the move m should not be included in the set M_n of moves after which enlivenment is possible. On the other hand, if $\eta_m = 1$, then the post-enlivenment node e_m^+ and ensuing subtree $T_{\geq e_m^+}^+$ should entirely replace the node n_m and ensuing subtree $T_{\geq n_m}$ in the subtree $T_{\geq n}^+$.

In any case, when enlivenment can occur only along the particular edge $n \rightarrow n_m$ emanating from the initial node n in the subtree $T_{\geq n}$, then the original continuation subtree $T_{\geq n_m}$ remains unenlivened. It follows that its enlivened evaluation $v^+(T_{\geq n_m})$ equals its unenlivened evaluation $v(T_{\geq n_m})$. So, applying the obvious counterpart of rule (56) for the particular event node e_m^- of $T_{\geq n}^+$ tells us that, in case $\eta_m \in (0, 1)$, one has

$$v^+(T_{\geq e_m^-}^+) = (1 - \eta_m) v(T_{\geq n_m}) + \eta_m v^+(T_{\geq e_m^+}^+) \quad (62)$$

Finally, in order to find the enlivened evaluation $v^+(T_{\geq n}^+)$ of the entire enlivened continuation subtree $T_{\geq n}^+$, we apply the obvious modification for this tree of whichever of the three rules (54), (56), and (58) is relevant at each node n' of $N_{\geq n}^+$, according as n' is a chance, event, or decision node.

8.7 Bayesian Rationality in Simply Enlivened Trees

Given any node $n \in N$, consider once again any simple enlivenment $T_{\geq n}^+$ of its continuation decision subtree $T_{\geq n}$ that starts at n with, for any move $m \in M_n^+ \subseteq M_n = N_{\geq n}^{+1}$, the stem $e_m^- \rightarrow e_m^+$ shown in Figure 7, as well as the continuation subtrees $T_{\geq e_m^-}^+$ and $T_{\geq e_m^+}^+$. Note that both these subtrees belong to the enlivened domain $\mathcal{T}^{S^+}(Y^+)$ of finite decision trees.

In effect, therefore, a simple enlivenment $T^+ = (N^+, E^+)$ of a decision tree $T = (N, E)$ is one in which multiple minimal enlivenments occur at all the buds e_m^- , where m is a potential move in a finite set $\cup_{n \in N \setminus N^t} M_n$. Each of these buds e_m^- gets inserted into an edge of T that belongs to the set

$$E = \cup_{n \in N \setminus N^t} \{n \rightarrow n_m \mid m \in M_n\}$$

So in principle an unboundedly rational agent, using what could be an unbounded model, could use backward recursion in order to complete the calculation of: (i) first, all the evaluations $\hat{v}^+(T_{\geq e_m^+}^+)$ of unenlivened continuation

subtrees in the simple enlivenment, in each of which the initial node is the unique post-enlivenment node e_m^+ that immediately follows a bud e_m^- with $m \in M_n$; (ii) second, all the enlivened evaluations $\hat{v}^+(T_{\geq n}^+)$ of the continuation subtrees $T_{\geq n}^+$ which start at any node n of the original unenlivened tree T .

8.8 Recursive Enlivenments

So far our discussion has been limited to simply enlivened decision trees. These are trees where any enlivenment edge $e_m^- \rightarrow e_m^+$ of the kind which was discussed in Section 8.4 must be inserted directly into an edge $n \rightarrow n_m$ of the original decision tree T , or of any its continuation subtrees $T_{\geq n}$, as illustrated in Figure 7. This requirement implies that, for any $m \in M_n$, the continuation subtree $T_{\geq e_m^+}^+$ starting after the enlivening event at e_m^- is always a decision tree in which no further enlivenment could occur.

Yet in the account of the Homeric example in Section 3, the decision tree that Kirke described to Odysseus was enlivened in several successive stages. Indeed, there were subsequent enlivenings of decision trees that had emerged only after previous enlivenings. This makes evident the possibility of *recursive enlivenment*. And of course anybody who has ever played Chess at any level beyond the most basic also knows that recursive enlivenment affects a player’s evolving understanding of the game being played, and so of how to evaluate any position in that game.

It should be evident by now, if it was not already, that my attempts to analyse enlivened decision trees in general have gone as far as they should in this paper. It remains to consider next the unavoidable truncations of enlivened trees.

9 Enlivened but Truncated Decision Trees

“The second factor which imposes a horizon upon the imaginative creation of the future is that uncertainty becomes more and more unbounded by considerations of what is possible, the more remote the date considered.” — Shackle (1969, page 224)

9.1 Extreme Truncation

Section 7 introduced the notion of a truncated decision tree as a natural form of bounded model for a decision tree that is too complicated for a full analysis to be practical. Applied to the reduction of a minimally enlivened continuation decision tree $T_{\geq n}^+$, whose influence diagram was shown in Section 8.1 as Figure 7, an extreme truncation removes the continuation subtree $T_{\geq e_m^+}^+$. This makes its initial node, which is the post-enlivenment node e_m^+ , into a truncation node which has become the terminal node of the agent's truncated tree $\hat{T}_{\geq n}^+$, whose influence diagram corresponding to one path is shown in Figure 8. In particular, the terminal node e_m^+ has become a terminal consequence node. The resulting consequence $\hat{\gamma}(e_m^+)$ is the real-valued subjective evaluation which should be based on some estimate of what the decision maker's expected utility at e_m^+ would have been the continuation subtree $T_{\geq e_m^+}^+$ had remained.




Figure 8: A Bud and Stem with a Subjective Evaluation

This form of extreme truncation obviously reduces the minimally enlivened decision tree T^+ to a truncated tree \hat{T}^+ in which any enlivenment stem $e_m^- \rightarrow e_m^+$ inserted into an edge $n \rightarrow n_m$ joins:

1. a bud e_m^- in the form of an event node at which there is a subjective probability $\eta_m \in (0, 1)$ of enlivenment occurring and taking the agent to node e_m^+ which ends the stem $e_m^- \rightarrow e_m^+$, instead of to the node n_m which ends the edge $n \rightarrow n_m$ in the original tree;
2. the node e_m^+ which becomes a truncation node that has an evaluation $\hat{\gamma}(e_m^+) \in \mathbb{R}$ attached.

But the same kind of reduction still works even in a recursively enlivened tree which is not simply enlivened because there exists a non-empty set of enlivened edges $n \rightarrow n_m$ in the original tree T at each of which the continuation subtree $T_{\geq e_m^+}^+$ after the enlivenment stem $e_m^- \rightarrow e_m^+$ of Figure 7 is subject to further enlivenment.

Now, given any enlivened continuation subtree $T_{\geq n}^+$ that results from such recursive enlivenment, provided that the set $N_{\geq n}^+$ of nodes is still finite, in principle one could still employ standard backward recursion in order to calculate the appropriate valuation attached to each continuation subtree $T_{\geq e_m^+}^+$ of $T_{\geq n}^+$.

9.2 Enlivened but Truncated Decision Trees

In many cases, however, it will be practically impossible to complete in full the computation by backward recursion of the evaluation $v(T_{\geq n})$ of every relevant continuation subtree $T_{\geq n}$ of the original tree T , let alone all the relevant evaluations of continuation subtrees in what may be the far more complicated corresponding computations for the whole of the enlivened tree T^+ . These computational difficulties make it practically imperative to truncate the agent's enlivened decision tree. Following the analysis in Section 7, this creates the need to attach subjective evaluations to at least some of the terminal nodes of a truncated decision tree. Indeed, for games as complicated as Chess or Go, the typical case requires the agent to attach to all but a very few terminal nodes a subjective evaluation rather than a consequence in the form of a definite result. This helps explain why some form of Monte Carlo tree search algorithm discussed in Section 7.5 has played such a key role in improving algorithms for playing Chess or Go.

Accordingly, following Section 7, consider any truncation that reduces some continuation decision trees to trivial trees consisting of a single truncation node with a real-valued evaluation as a consequence. Ideally, even if there is recursive enlivenment, those evaluations should be what would emerge from a complete backward recursive construction applied to the whole enlivened decision tree, without any truncation. Using the subjective evaluations $\hat{\gamma}(e_m^+)$ in the extremely truncated tree, however, obviates entirely any need for any recursive calculation. Of course, though not strictly necessary, some recursive calculations could help to produce normatively superior beliefs about what subjective evaluations should be attached to the continuation subtrees $T_{\geq e_m^+}^+$ emanating from the relevant post-enlivenment nodes e_m^+ of the enlivened decision tree T^+ .

The required extension to truncated decision trees with subjective evaluations replaces the relevant domain $L^{S^+}(Y^+)$ of possible enlivened AA consequence lotteries that was defined in Section 8.5 with the new enlivened

AA consequence lottery domain $L^{S^+}(Y^+ \cup \mathbb{R})$. We also extend the earlier domain $\mathcal{T}^{S^+}(Y^+)$ of untruncated enlivened decision trees in order to accommodate the range of all possible subjective evaluations $\hat{v}(n)$ of AA consequence lotteries $L^{S^+}(Y^+ \cup \mathbb{R})$ that could be attached to a truncation node of a decision tree. This range of possible evaluations consists of the expected value for any lottery in $L^{S^+}(Y^+ \cup \mathbb{R})$ of any Bernoulli utility function $Y^+ \ni y \mapsto u^+(y) \in \mathbb{R}$ that has been normalized to satisfy the two restrictions $u^+(\underline{y}) = \underline{u}$ and $u^+(\bar{y}) = \bar{u}$ of (51), as described in Section 6.5. Because of these normalization restrictions, the relevant extension of the domain $\mathcal{T}^{S^+}(Y^+)$ of enlivened decision trees depends on the four parameters $\underline{y}, \underline{u}, \bar{y}, \bar{u}$ that together belong to the Cartesian product space $Y \times \mathbb{R} \times Y \times \mathbb{R}$. Accordingly, in order to recognize this dependence, we let $\hat{\mathcal{T}}^{S^+}(Y^+, \underline{y}, \underline{u}, \bar{y}, \bar{u})$ denote the domain of enlivened but truncated finite decision trees, including the evaluations of their truncation nodes.

Having extended the consequence domain in this way, the next step is to extend the construction in Section 6.7 of the evaluation $v(T_{\geq n})$ of each continuation subtree $T_{\geq n}$ of each decision tree $T \in \mathcal{T}^S(Y^S)$. Let $\hat{v}^+(\hat{T}_{\geq n}^+)$ denote the result of the extended construction, which is an evaluation defined for each continuation subtree $\hat{T}_{\geq n}^+$ of each tree \hat{T}^+ in the domain $\hat{\mathcal{T}}^{S^+}(Y^+, \underline{y}, \underline{u}, \bar{y}, \bar{u})$ of enlivened but truncated decision trees. These evaluations should still satisfy at all nodes whichever of the four recurrence relations (53), (54), (56), and (58) is relevant, though with each $v(T_{\geq n})$ and each $v(T_{\geq n'})$ replaced by $\hat{v}^+(\hat{T}_{\geq n}^+)$ and $\hat{v}^+(\hat{T}_{\geq n'}^+)$ respectively. The only new feature is that, at any terminal evaluation node n with normalized subjective evaluation $\hat{\gamma}(n)$, the equation (53) should be replaced by the obvious

$$\hat{v}^+(\hat{T}_{\geq n}^+) = \hat{\gamma}(n) \tag{63}$$

9.3 Characterizing Bayesian Rationality

Section 1.1 offered a summary of the results in Hammond (1988a, b; 1998a, b; 1999, 2022) that characterize Bayesian rational decision-making. The main conclusion of this paper is that these results survive, though with normalized subjective evaluations at truncation nodes, for the domain $\hat{\mathcal{T}}^{S^+}(Y^+, \underline{y}, \underline{u}, \bar{y}, \bar{u})$ of finite enlivened but truncated decision trees. These results all follow from adapting in an obvious way the earlier results for the old domain of unenlivened and untruncated finite decision trees to the new domain of enlivened but truncated finite decision trees. This is because there is an obvious em-

bedding of the new domain in the old domain, for which previous work had characterized Bayesian rationality.

9.4 Limitations of Bayesian Rationality

A concept of rationality that is more philosophically refined than mere Bayesian rationality would presumably require a rational agent to use some kind of some kind of “normatively justified” Bernoulli utility function defined on the consequence domain, together with some kind of “normatively justified” subjective probabilities over uncertain states of the world. Thus, this richer concept of rationality would go beyond mere Bayesian rationality.

Bayesian rationality in enlivened trees is no less limited. Indeed, full rationality should require the agent’s subjective evaluations at terminal evaluation nodes to be normatively justified, in addition to the agent’s normalized Bernoulli utility function and subjective probabilities. Of course, the agent’s estimates of the relevant subjective probabilities and subjective evaluations may well be improved by procedures such as Monte Carlo tree search, as considered in Section 7.5.

An alternative procedure that has proved useful in some business applications involves “scenario planning”. This is somewhat similar to the rational shortlist method considered in Section 2.2. It involves detailed consideration of a range of scenarios that are chosen in an attempt to cover eventualities that are likely to have a significant bearing on the choice of policy. This is in contrast to what may be a standard view that the focus should be on the most likely eventualities.²⁰

10 Extensions and Conclusions

10.1 Decision Trees with Timed Consequence Nodes

In Hammond and Troccoli Moretti (2025) we consider finite decision trees which, in addition to decision, chance, event, and terminal nodes, also have timed consequence nodes. Given the consequence domain Y , for each $m \in \mathbb{N}$

²⁰See, for example, Jefferson (2014) for the point of view of Shell’s former Chief Economist, and his experience with scenario planning. See also Derbyshire (2017) for discussion of the possible links between scenario planning and Shackle’s concept of potential surprise.

let Y^m denote the Cartesian product of m copies of Y . Then let $\mathbf{Y}(Y)$ denote the domain of all possible timed consequence streams which, for some $m \in \mathbb{N}$, take the form $\mathbf{y} = (t_j, y_j)_{j=1}^m \in \mathbb{R}^m \times Y^m$, where $t_1 < t_2 < \dots < t_m$. With this construction, any deterministic decision tree T with timed consequence nodes can be reduced to an *equivalent* decision tree $T^* \in \mathcal{T}(\mathbf{Y}(Y))$ in which:

1. at each terminal node $n \in N^t$, which is a consequence node, the assigned terminal consequence $\gamma(n) \in Y$ is replaced by the unique timed consequence stream $\mathbf{y}(n) \in \mathbf{Y}(Y)$ that ends at node n ;
2. each timed consequence node is then eliminated.

In addition to deterministic decision trees, a similar reduction is possible for all decision trees that have chance and/or event nodes.

In a decision tree with timed consequence nodes, the eventual consequence of reaching a terminal node n is, not a lottery over single consequences at node n , but rather a lottery over the stream of consequences and consequence lotteries that accumulate along the unique path through the decision tree that ends at node n . The implications of allowing timed consequence nodes are then routine unless the consequence of reaching any consequence node n includes a “menu consequence” which depends on what consequence lotteries are feasible in the continuation decision tree emanating from node n . Even then, however, the results summarized in Section 1.1 that justify Bayesian rationality remain valid; all that changes is that the domain of relevant consequence streams becomes much richer. These changes can be significant, however. For example, they allow scope for preferences for flexibility, as in Koopmans (1950, 1964) and Kreps (1979). They also allow preferences for the timing of resolution of uncertainty, as in Kreps and Porteus (1978). Also possible are menu consequences that reflect preferences over lotteries such as those considered by Machina (1989) and by Gilboa and Schmeidler (1989) which cannot be represented by an expected utility function, as well as recursive preferences for intertemporal consumption streams of the kind that Epstein and Zin (1989) first considered.

Indeed, for trees with timed consequence nodes as well as decision, chance and event nodes, there is an obvious extension of the rules set out in the four equations (53), (54), (56), and (58) of Section 6.7. This extension treats the case when node n is a consequence node, in which case the timed consequence attached to node n is prepended at the beginning of each consequence stream \mathbf{y} that arises in the continuation subtree $T_{\geq n}$. With this extension,

the characterization in Section 9.3 of Bayesian rationality for enlivened but untruncated decision trees still applies.

A minor complication does arise, however, in characterizing Bayesian rationality with subjective evaluations in enlivened but truncated decision trees. The issue is how to truncate a non-trivial stream of consequences and consequence lotteries that accumulates along a path through the decision tree. The obvious remedy is simply to append as a consequence attached to the truncation node a real-valued subjective evaluation of the kind considered in Sections 7.3, 9.2, and 9.3 to the consequence stream that has accumulated before the truncation.

10.2 Reverse Bayesianism

“Reverse Bayesianism” was described in the series of joint papers by Karni and Vierø, as well as those by Vierø on her own, that were cited in Section 1.3. For the general finite decision trees considered here, reverse Bayesianism is the result saying that, if you condition the probabilities of different consequences in any enlivened decision tree T^+ on the event that enlivening does not occur, the result should be the corresponding probabilities in the original unenlivened decision tree T .

10.3 Transformative Experiences

The concept of transformative experiences arose in philosophy thanks to Paul (2014, 2015a, b, c), though similar ideas were discussed earlier in Ullmann-Margalit’s (2006) paper on “big decisions” whose first characteristic (p. 158) is that they must be “transformative, or ‘core affecting’”. As Paul (2015b, p. 761) writes:

a transformative experience is . . . both radically new to the agent and changes her in a deep and fundamental way; there are experiences such as becoming a parent, discovering a new faith, emigrating to a new country, or fighting in a war. . . .

An *epistemically* transformative experience is an experience that teaches you something you could not have learned without having that kind of experience. Having that experience gives you new abilities to imagine, recognize, and cognitively model possible future experiences of that kind. A *personally* transformative

experience changes you in some deep and personally fundamental way, for example, by changing your core personal preferences or by changing the way you understand your desires and the kind of person you take yourself to be. A *transformative experience*, then, is an experience that is both epistemically and personally transformative.

The main problem with transformative decisions is that our standard decision models break down when we lack epistemic access to the subjective values for our possible outcomes.

As two philosophically prominent examples of transformative experiences, she considers the decisions of whether to become a vampire (!) or to have a child.²¹ See also the discussion by Pettigrew (2015), Barnes (2015), and Campbell (2015), as well as later by writers such as Bykvist and Stefánsson (2017) and Randell (2023). The claim we make here is that any transformative experience involves a decision tree that becomes enlivened in case it includes any decision by the agent which determines whether to undergo the novel experience being considered.

10.4 A Possible Conclusion

The widely quoted aphorism due to the statistician George Box that was quoted at the head of Section 1.2 should remind us of the inevitable limitations that will occur in any formal model of all but the most trivial decision problems. Indeed, examples that extend in time from Homer's *Odyssey* to modern algorithms for playing chess demonstrate that, for an agent who has one or more decisions to make, the usefulness of any model is all too likely to be temporary. This paper has begun an investigation of what decisions are rational for an agent who recognizes this fundamental difficulty. Specifically, it is argued that the best which such an agent can do, in effect, is to construct a subjective truncated probabilistic model of what the ultimate ex post evaluation of each possible decision could be, and then to maximize the expectation of this evaluation.

²¹Perhaps it would make a bit more sense to ask if it would be better to start practising vampirism, or occasionally adopting the feeding habits of a vampire bat.

Acknowledgements: The research reported here was supported from 2007 to 2010 by a Marie Curie Excellence Chair funded by the European Commission under contract number MEXC-CT-2006-041121. Many thanks also to Kenneth Arrow, Ken Binmore, and Burkhard Schipper for their helpful suggestions during these early years, to Stan Metcalfe for increasing my appreciation of Shackle’s original ideas, as well as to Marcus Miller and Joanne Yoong for enlivened discussions, while absolving them of all responsibility for my errors or omissions. Many thanks also for their patient attention and encouragement to audiences at: the GSB/Economics Department theory seminar at Stanford University (March 2007); the London School of Economics conference on preference change (May 2009); the University of Warwick (November 2010); CORE at the Université Catholique de Louvain (September 2011); the Hausdorff Center for Mathematics in Bonn (August 2013); the Università Cattolica del Sacro Cuore in Milan (October 2013); the conference on unawareness in game theory at the University of Queensland (February 2014); the conference on game theory and applications at the University of Bristol (May 2016); the Microeconomics Work in Progress online seminar at the University of Warwick (March 2021); the workshop on “Issues in Dynamic Decision Theory” at the University of Konstanz (July 2023); the Centre d’Économie de la Sorbonne in Paris (November 2023); and the IX Hurwicz Workshop on Mechanism Design Theory at the SGH Warsaw School of Economics and at the Banach Center of the Institute of Mathematics of the Polish Academy of Sciences (December 2024). Recently, the prolonged process of producing the current version of this paper was greatly helped by discussions with Agustín Troccoli Moretti, my coauthor on a related project, as well as with my colleagues Kirk Surgener and Pablo Beker at Warwick. Finally, this greatly improved version of an earlier CRETA working paper owes much to the helpful comments of some unusually diligent anonymous referees.

References

- Abramson, B. (1991) *The Expected-Outcome Model of Two-Player Games* (Morgan Kaufman).
- Anscombe, F.J., and R.J. Aumann (1963) “A Definition of Subjective Probability” *Annals of Mathematical Statistics* 34: 199–205.
- Barnes, E. (2015) “What You Can Expect When You Don’t Want to be Expecting” *Philosophy and Phenomenological Research* 91(3): 775–786.

- Browne, C.B. *et al.* (2012) “A Survey of Monte Carlo Tree Search Methods” *IEEE Transactions on Computational Intelligence and AI in Games* 4 (1): 1–43.
- Bykvist, K., and H.O. Stefánsson (2017) “Epistemic Transformation and Rational Choice” *Economics and Philosophy* 33(1): 125–138.
- Campbell, J. (2015) “L.A. Paul’s *Transformative Experience*” *Philosophy and Phenomenological Research* 91(3): 787–793.
- Dekel, E., B.L. Lipman, and A. Rustichini (2001) “Representing Preferences with a Unique Subjective State Space,” *Econometrica* 69: 891–934.
- Dekel, E., B.L. Lipman, A. Rustichini, and T. Sarver (2007) “Representing Preferences with a Unique Subjective State Space: Corrigendum” *Econometrica* 75: 591–600.
- Derbyshire, J. (2017) “Potential Surprise Theory as a Theoretical Foundation for Scenario Planning” *Technological Forecasting and Social Change* 124: 77–87.
- Drèze, J.H. (1962) “L’utilité sociale d’une vie humaine” *Revue Française de Recherche Operationelle* 23: 93–118,
- Drèze, J.H., and A. Rustichini (2004) “State-Dependent Utility and Decision Theory” in S. Barberà, P.J. Hammond and C. Seidl (eds.) *Handbook of Utility Theory, Vol. 2: Extensions* (Dordrecht: Kluwer Academic) ch. 16, pp. 839–892.
- Earl, P., and B. Littleboy (2014) *G.L.S. Shackle* (Basingstoke: Palgrave Macmillan).
- Elster, J. (1979) *Ulysses and the Sirens: Studies in Rationality and Irrationality* (Cambridge: Cambridge University Press).
- Epstein, L.G., and S.E. Zin (1989) “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework” *Econometrica* 57 (4): 937–969.
- Gilboa, I., and D. Schmeidler (1989) “Maxmin Expected Utility with Non-Unique Prior” *Journal of Mathematical Economics* 18: 141–153.
- Grant, S. and J. Quiggin (2015a) “A Preference Model for Choice Subject to Surprise” *Theory and Decision* 79 (1): 167–180.

- Grant, S., J. Kline, P. O’Callaghan, and J. Quiggin (2015b) “Sub-models for Interactive Unawareness” *Theory and Decision* 79 (4): 601–613.
- Halpern, J.Y. and L.C. Rêgo (2014) “Extensive Games with Possibly Unaware Players” *Mathematical Social Sciences* 70: 42–58.
- Hammond, P.J. (1976) “Changing Tastes and Coherent Dynamic Choice” *Review of Economic Studies* 43: 159–173.
- Hammond, P.J. (1988a) “Consequentialism and the Independence Axiom” in B.R. Munier (ed.) *Risk, Decision and Rationality (Proceedings of the 3rd International Conference on the Foundations and Applications of Utility, Risk and Decision Theories)* (Dordrecht: D. Reidel), pp. 503–516.
- Hammond, P.J. (1988b) “Consequentialist Foundations for Expected Utility” *Theory and Decision* 25: 25–78.
- Hammond, P.J. (1998a) “Objective Expected Utility: A Consequentialist Perspective” in S. Barberà, P.J. Hammond and C. Seidl (eds.) *Handbook of Utility Theory, Vol. 1: Principles* (Dordrecht: Kluwer Academic) ch. 5, pp. 145–211.
- Hammond, P.J. (1998b) “Subjective Expected Utility” in S. Barberà, P.J. Hammond and C. Seidl (eds.) *Handbook of Utility Theory, Vol. 1: Principles* (Dordrecht: Kluwer Academic) ch. 6, pp. 213–271.
- Hammond, P.J. (1999) “Subjectively Expected State-Independent Utility on State-Dependent Consequence Domains” in M.J. Machina and B. Munier (eds.) *Beliefs, Interactions, and Preferences in Decision Making* (Dordrecht: Kluwer Academic), pp. 7–21.
- Hammond, P.J. (2007) “Schumpeterian Innovation in Modelling Decisions, Games, and Economic Behaviour” *History of Economic Ideas* XV: 179–195.
- Hammond, P.J. (2022) “Prerationality as Avoiding Predictably Regrettable Consequences” *Revue Économique* 73: 943–976.
- Hammond, P.J. and A. Troccoli Moretti (2025) University of Warwick, CRETA working paper (in preparation).
- Hansen, L.P. and T. Sargent (2007) *Robustness* (Princeton: Princeton University Press).

- Hicks, J.R. (1946) *Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory* (Oxford: Clarendon Press).
- Howard, R.A. and J.E. Matheson (2005) “Influence Diagram Retrospective” *ACM Decision Analysis* Vol. 2, No. 3. <https://doi.org/10.1287/deca.1050.0050>
- Jefferson, M. (2014) “The Passage of Time: Shackle, Shell and Scenarios” in Earl, P., and B. Littleboy (eds.) *G.L.S. Shackle* (Basingstoke: Palgrave Macmillan) pp. 198–214.
- Karni, E. (1985) *Decision Making Under Uncertainty: The Case of State-Dependent Preference* (Cambridge, MA: Harvard University Press).
- Karni, E. and M.L. Vierø (2013) “‘Reverse Bayesianism’: A Choice-Based Theory of Growing Awareness” *American Economic Review* 103, 2790–2810.
- Karni, E. and M.L. Vierø (2015) “Probabilistic Sophistication and Reverse Bayesianism” *Journal of Risk and Uncertainty* 50: 189–208.
- Karni, E. and M.L. Vierø (2017) “Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance” *Journal of Economic Theory* 168: 301–328.
- Koopmans, T.C. (1950) “Utility Analysis of Decisions Affecting Future Well-Being” (abstract) *Econometrica* 18: 174–5.
- Koopmans, T.C. (1964) “On Flexibility of Future Preference” in M.W. Shelly and G.L. Bryan (eds.) *Human Judgments and Optimality* (New York: John Wiley), ch. 13, pp. 243–254.
- Kreps, D.M. (1990) *Game Theory and Economic Modelling* (Oxford: Clarendon Press).
- Kreps, D.M. (1992) “Static Choice in the Presence of Unforeseen Contingencies” in P. Dasgupta, D. Gale, O. Hart, and E. Maskin (eds.) *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn* (Cambridge, Mass.: M.I.T. Press), pp. 258–281.
- Kreps, D.M., and E.L. Porteus (1978) “Temporal Resolution of Uncertainty and Dynamic Choice Theory” *Econometrica* 46 (1): 185–200.
- Machina, M. (1989) “Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty” *Journal of Economic Literature* 27: 1622–1668.

- Manzini, P. and M. Mariotti (2007) “Sequentially Rationalizable Choice” *American Economic Review* 97: 1824–1839.
- Marschak, J. (1950) “Rational Behavior, Uncertain Prospects, and Measurable Utility” *Econometrica* 18 (2): 111–141.
- Metcalf, S., S. Salles-Filho, L.T. Duarte, A. Bin, A.T. Azevedo, and P.H.A. Feitosa (2021) “Shackle’s Approach Towards Priority Setting and Decision-Making in Science, Technology, and Innovation” *Futures* 134: 102838.
- Myerson, R.B. (1979) “An Axiomatic Derivation of Subjective Probability, Utility, and Evaluation Functions” *Theory and Decision* 11: 339–352.
- Myerson, R.B. (1983) “A Dynamic Microeconomic Model with Durable Goods and Adaptive Expectation” *Journal of Economic Behaviour and Organization* 4: 309–351.
- Paul, L.A. (2014) *Transformative Experience* (Oxford U. Press).
- Paul, L.A. (2015a) “What You Can’t Expect When You’re Expecting” *Res Philosophica* [online] 92 (2): 1–22.
- Paul, L.A. (2015b) “Précis of *Transformative Experience*” *Philosophy and Phenomenological Research* 91 (3): 760–764.
- Paul, L.A. (2015c) “Transformative Experience: Replies to Pettigrew, Barnes and Campbell” *Philosophy and Phenomenological Research* 91 (3): 794–813.
- Pettigrew, R. (2015) “Transformative Experience and Decision Theory” *Philosophy and Phenomenological Research* 91 (3): 766–774.
- Pollak, R.A. (1968) “Consistent Planning” *Review of Economic Studies* 35: 201–208.
- Raiffa, H. (1968) *Decision Analysis: Introductory Lectures on Choices under Uncertainty* (Addison-Wesley).
- Randell, P. (2023) “Familiar Transformative Experiences” *Synthese* 202: 45.
- Savage, L.J. (1954, 1972) *Foundations of Statistics* (New York: John Wiley; and New York: Dover Publications).
- Schipper, B.C. (2014a) “Unawareness—A Gentle Introduction to Both the Literature and the Special Issue” *Mathematical Social Sciences* 70: 1–9.

- Schipper, B.C. (2014b) “Preference-Based Unawareness” *Mathematical Social Sciences* 70: 34–41.
- Schumpeter J.A. (1911) [2nd edn. 1926] *Theorie der wirtschaftlichen Entwicklung; Eine Untersuchung über Unternehmergewinn, Kapital, Kredit, Zins und den Konjunkturzyklus* (Leipzig, Duncker & Humblot).
- Schumpeter J.A. (1934) [1961] *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle* translated from Schumpeter (1911) by R. Opie, with a new introduction by J.E. Elliott.
- Schervish, M.J., T. Seidenfeld, and J.B. Kadane (1990) “State-Dependent Utilities” *Journal of the American Statistical Association* 85 (411): 840–847.
- Seidenfeld, T., M.J. Schervish, and J.B. Kadane (2010) “Coherent Choice Functions under Uncertainty” *Synthese* 172: 157–176.
- Shackle, G.L.S. (1953) “The Logic of Surprise” *Economica, New Series* 20: 112–117.
- Shackle, G.L.S. (1969) *Decision, order, and time in human affairs* (Cambridge: Cambridge University Press).
- Silver, D. *et al.* (2018) “A General Reinforcement Learning Algorithm that Masters Chess, Shogi, and Go Through Self-Play” *Science* 362 (issue 6419): 1140–1144.
- Simon, H.A. (1955) “A Behavioral Model of Rational Choice” *Quarterly Journal of Economics* 69: 99–118.
- Simon, H.A. (1957) *Models of Man* (New York: John Wiley).
- Sims, C.A. (2003) “Implications of Rational Inattention” *Journal of Monetary Economics* 50 (3): 665–690.
- Sims, C.A. (2011) “Rational Inattention and Monetary Economics” *Handbook of Monetary Economics, vol. 3* (North-Holland) pp. 155–181.
- Strotz, R.H. (1956) “Myopia and Inconsistency in Dynamic Utility Maximization” *Review of Economic Studies* 23: 165–180.
- Ullmann-Margalit, E. (2006) “Big Decisions: Opting, Converting, Drifting” *Royal Institute of Philosophy Supplement* 58: 157–172.

- Vierø, M.-L. (2009) “Exactly What Happens After the Anscombe–Aumann Race? — Representing Preferences in Vague Environments” *Economic Theory* 41, 175–212.
- Vierø, M.-L. (2021) “An Intertemporal Model of Growing Awareness” *Journal of Economic Theory* 197: 105351.
- Von Neumann, J. (1928) “Zur Theorie der Gesellschaftsspiele” *Mathematische Annalen* 100: 295–320.
- Zermelo, E. (1913) “Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels” *Proceedings of the Fifth International Congress of Mathematicians*, Vol. II, pp. 501–504; republished with English translation as “On an Application of Set Theory to the Theory of the Game of Chess” in E. Zermelo *et al.* (2010) *Collected Works/Gesammelte Werke: Volume I/Band I — Set Theory, Miscellanea/Mengenlehre, Varia: 1* (Heidelberg: Springer).