Rationality in Finite Decision Trees with Consequence Chains That may Include Menus*

Peter J. Hammond† Agustín Troccoli Moretti‡

University of Warwick UPF and BSE

May 17, 2026
typeset from `menusSITE2026a.tex`

Abstract

Following the work in Hammond and Troccoli Moretti (2026) on rational behaviour in finite decision trees with consequence chains, we focus here on consequences that may include a "menu" of consequence chains that are feasible in the continuation subtree which starts at the relevant consequence node. Introducing menu consequences allows prerational behaviour rules, as defined in Hammond (2022), to accommodate a plethora of phenomena that otherwise superficially violate "consequentialism". Prominent examples include apparent violations of the ordinality and independence axioms of expected utility theory. Other instances of "menu effects" arise because of preferences for when uncertainty is resolved, as well as Epstein–Zin (1989) preferences, and the influence of temptation.

Keywords: *prerationality, consequentialist decision theory, decision trees, consequence nodes, normal form invariance, menu consequences*

JEL Classification: D81, D91, D11, D63, E21, E71, G41

*Earlier versions of this paper were included in Chapter 2 of Agustín's PhD dissertation. They were also presented to the Economic Theory workshop at Warwick, the SAET 2024 conference in Manchester, the IX Hurwicz Workshop in Warsaw in 2024, and to a lecture at the Centre d'Économie de la Sorbonne in Paris. For their helpful discussion we would like to thank Agustín's thesis adviser Herakles Polemarchakis and his thesis examiners Pablo Beker and Georgios Gerasimou, as well as audience members at the different seminars, including Marcus Pivato and Raphael Oliveira Gomes.

†CRETA, Department of Economics, University of Warwick, Coventry CV4 7AL, UK. E-mail: p.j.hammond@warwick.ac.uk

‡Departament d'Economia i Empresa, UPF, Jaume I Building, Ramon Trias Fargas, 25-27, 08005 Barcelona and Barcelona School of Economics. E-mail: agustin.troccoli@upf.edu

1 Introduction and Outline

1.1 Menus and Consequentialization

In Hammond and Troccoli Moretti (2026), which we refer to from now on as H&TM, we have developed a consequentialist decision theory for finite decision trees with chains of *timed consequence nodes*, of which only the last can be a terminal node. The present contribution extends that framework by allowing any intermediate timed consequence to include a *menu consequence* which records, as part of the consequence chain, the set of continuation prospects that are feasible in the relevant continuation subtree. This extension leaves unchanged the basic primitives and the behavioural notion of prerationality. But it does substantially expand the scope of what can be represented as consequentialist behaviour in dynamic choice problems.

Menus have long been central objects of study in decision theory. In deterministic environments beginning with Koopmans (1950, 1964), they were used by Kreps (1979) and others to formalise any preference for flexibility and related phenomena. In risky environments, once the domain was enlarged to allow lotteries as objects of choice, the menus framework became a workhorse for modelling psychological choice phenomena, as in the seminal contributions of Dekel et al. (2001, 2007) and of Gul and Pesendorfer (2001). A large subsequent literature has used preferences over menus to formalise, amongst other phenomena, anticipated regret (Sarver, 2008), self-deception (Kopoylov and Noor, 2018), as well as shame and guilt (Dillenberger and Sadowski, 2012; Noor and Ren, 2003). See Lipman and Pesendorfer (2013) for a survey.

The literature on menus also connects rather naturally to the decision-theoretic rationale for the consequentialist doctrine set out in Hammond (1988). There it is stressed that consequentialism constrains *how* behaviour is to be justified, not *which* features of a decision problem may enter the justification. In particular, to quote from p. 26:

> *"[The doctrine] would be false if missed opportunities, regrets, sunk costs, etc., affected behaviour and yet were excluded from the domain of consequences. As a normative principle, however, consequentialism requires everything which should be allowed to affect decisions to count as a relevant consequence — behaviour is evaluated by its consequences, and nothing else. If regrets, sunk costs, even the structure of the decision tree itself are relevant to normative behaviour, they are therefore already in the consequence domain."*

This observation highlights the importance of incorporating every aspect of the decision problem that is relevant to the decision maker's welfare, whether it be regrets, tempting alternatives, or sunk costs, into the description of a consequence. Some philosophers, especially Portmore (2007, 2022, 2023) and Campbell (2011), refer to this process as *consequentialisation*. It even allows deontological concepts from Kantian ethics and from Harsanyi's rule utilitarianism to be included.

A central implication of these ideas is that whether a given behavioural pattern appears consequentialist depends on how the consequence domain is specified. Apparent non-consequentialist behaviour may therefore be an artefact of an incompletely specified consequence domain, in the sense that the same behaviour would satisfy the requirement of consequentialism once the domain has been enriched appropriately. The discussion in Hammond (1988), however, leaves open the question of how to implement this prescription formally in a general and tractable way.

1.2 Introducing Menu Consequences

Conceptually, our paper is an exploration in the possibilities created by consequentializing menus in the form of a set of more basic consequences. This accords with the consequentialist doctrine which emphasises that whatever is normatively relevant for the decision maker should be incorporated into the description of a consequence. It parallels the requirement by Savage (1954) that each state be a complete description of all payoff-relevant contingencies, so that conditioning on the state removes all residual uncertainty. Here the analogue is that a consequence should be described richly enough to allow to encode any aspect of the decision problem that the decision maker treats as payoff-relevant.

A key virtue of our approach is that this enrichment does *not* depend on the specific application. We do not have to tailor the primitive consequence space anew for each behavioural irregularity. Instead, by admitting menu consequences in a uniform way, we can represent within a single fixed domain many phenomena that are often labelled non-consequentialist precisely because they involve sensitivity to the structure of the tree or to foregone, unreached, or avoided branches. In our framework, these features enter through menu consequences, and therefore become part of the description of the consequences themselves.

The present paper applies the idea of consequentializing menus to a domain of finite decision trees. Our approach is to enrich the timed-consequence domain not by tailoring the primitive consequence space case by case, but by admitting a uniform new sort of intermediate consequence — namely, a *menu consequence* that records what is relevant about the set of consequences that are feasible in a relevant continuation subtree. The central modelling restriction is that, in a consequentialist analysis, not every syntactic detail of a continuation subtree should matter directly. What matters, instead, is its *consequentialist content*, namely the menu of feasible continuation consequences it generates. Accordingly, a menu consequence node introduces no branching and adds no new continuation opportunities. Rather, it merely attaches to the relevant menu consequence node n a timed consequence in the form of the set $F(T_{\geq n})$ of consequence chains that are feasible in the continuation subtree starting at n.

1.3 Consequentialist normal-form invariance

The behavioural primitives we adopt are exactly the same as in H&TM. A *behaviour rule* β specifies, at each decision node n of each tree in the domain we consider, and conditional on each possible history $h \in H$ of timed consequences that occur prior to n, a non-empty set of admissible moves. *Consequentialist normal-form invariance*, as defined in Section 2, requires that admissibility depend only on the feasible set $F(T)$ of lotteries over consequence chains generated by the tree, and not on the particular extensive-form presentation of that feasible set. This invariance principle implies that behaviour can be represented by a choice function defined on the domain of finite feasible sets of lotteries over consequence chains, and that it induces a base preference relation defined on the relevant lottery domain. A base preference relation is *prerational* if it is generated in this way by some consequentialist normal-form invariant behaviour rule.

Conceptually, our framework puts consequentialisation into practice in a way that does not depend on the specific application we are considering. By admitting menu consequences uniformly, we obtain a single fixed consequence domain in which many behavioural patterns that are often labelled non-consequentialist become compatible with prerationality because they are reinterpreted as sensitivity to components of the realised stream. In particular, when a menu consequence M appears on the path, the Bernoulli utility index attached to M may depend on a function of the continuation feasible set. So the structure of unreached or avoided branches can matter *through* the consequences themselves.

1.4 Outline

Section 2 recalls the key concepts introduced in H&TM, for decision trees T with timed consequences but without menu consequences. These concepts include, for each history h in the form of a chain of timed consequences, and for each tree T whose initial node is a decision node in the relevant domain: (i) the induced feasible set $F(T)$ of consequence chains; (ii) the history-dependent behaviour set $\beta_h(T)$ not only at the initial node of T, when it is a decision node, but also at the initial decision node of each continuation subtree of the tree T that starts with a decision node; (iii) consequentialist normal-form invariance of the behaviour rule; (iv) the induced history-dependent base preference relation $\succsim_h$ defined on lotteries over consequence chains.

Thereafter the next four Sections 3–6 each use a different kind of example in order to illustrate the flexibility that menu consequences allow. Specifically, in Section 3, menu consequences in a deterministic decision tree are used to describe: (i) the work by Koopmans (1950, 1964) and by Kreps (1979) on a preference for flexibility; (ii) a polite refusal to take the largest piece of cake, or the most comfortable chair; (iii) a potential addict's preference for inflexibility; (iv) an individual's preference for the freedom to choose, or for individual rights; (v) Sen's concepts of individual agency, as well as of capabilities and "functionings"; (vi) an example ascribed to Arrow which violates a version of Chernoff's (1954) contraction consistency axiom, which Sen (1970a) calls "property α".

In the risky decision trees considered in Section 4, menu consequences allow an expected-utility preference to represent arbitrary attitudes to the timing of information resolution, in the spirit of the critique by Kreps and Porteus (1978) of the payoff-vector approach — see Section 4.3. As shown in Section 4.2, the same idea also yields a transparent consequentialist embedding of apparently non-expected-utility choice rules over roulette lotteries, by recording the chosen roulette lottery as an intermediate singleton menu consequence, as in Equation (6). Moreover, the two-period Epstein–Zin preferences exhibited in Equation (8) are a prominent special case. Also, as discussed in connection with Equation (9) of Section 4.5, our framework reconciles the literature on the effects of menus with the consequentialist foundations for expected utility. It does so by showing how, without abandoning the independence axiom for the enlarged consequence domain, temptation and self-control phenomena can be explained by the dependence of welfare on the realised feasible set.

Section 5 presents two examples of menu consequences involving decision trees with uncertainty represented by event nodes, where a horse lottery is resolved. The first of these examples concerns the Ellsberg (1951) paradox, where risk aversion is distinguished from uncertainty aversion. The second example concerns the precautionary principle.

The final examples, which are from social choice theory, appear in Section 6. These start with "Machina's Mom", followed by Diamond's (1967) widely cited criticism of Harsanyi's (1953, 1955) justification of utilitarianism. The last two examples are of menu consequences intended to capture Rawls' maximin concept of justice, or prioritarianism. The latter in particular represents an attempt by some philosophers to have a more egalitarian ethical objective than utilitarianism allegedly provides.

The main theoretic results of the paper start to appear in Section 7, which extends the formal theory introduced in H&TM and summarized in Section 2 so that consequence chains can accommodate menu consequences. It explains why modelling a menu consequence as a set of feasible consequence chains in a continuation decision tree is the natural analogue of consequentialization. It also explains why allowing multiple menu consequences on a path forces one to consider a recursive construction of the relevant consequence domain involving the length of a consequence chain. Section 7.2 formalises construction of the enriched domains $(\mathbf{S}_\infty, \mathbf{M}_\infty)$ and the resulting domain of decision trees with timed menu consequences. Appendix A provides the closure, existence, and minimality arguments for this construction, culminating in Propositions A.1 and A.2.

Section 8 establishes the analogue of the main representation theorem in HTS. Indeed, Theorem 8.1 states that a base preference relation on the enriched domain of lotteries over consequence chains is prerational and mixture continuous if and only if it admits an expected-utility representation, as specified in Equation (14) . Thus, even on the enriched domain that admits menu consequences, prerationality retains its tight logical link with Bayesian rationality.

Section 9 offers a concluding summary.

2 Preliminaries

This section recalls the basic definitions and notation of the consequentialist decision theory with timed consequences developed in H&TM, which in turn extends the earlier discussion of consequentialist behaviour and objective expected utility set out in Hammond (1988, 1998, 2022). We keep the exposition self-contained but deliberately brief. For a fuller, detailed treatment of the theory, we refer the reader to H&TM.

2.1 Timed consequences, consequence chains, and histories

Let Y be a non-empty set of elementary consequences. A *timed consequence* is a pair $(t, y) \in \mathbb{R} \times Y$, where t indicates the time at which the consequence y is experienced.

A timed consequence chain of length $\tau \in \mathbb{N}$ is a sequence

$$\mathbf{y} = (\boldsymbol{y}_j)_{j=1}^{\tau} = \big((t_1, y_1), \ldots, (t_\tau, y_\tau)\big) \in (\mathbb{R} \times Y)^\tau \quad \text{with} \quad t_{j+1} > t_j \ \text{ for all } j < \tau$$

of τ successive timed consequences. The set $\cup_{\tau \in \mathbb{N}} (\mathbb{R} \times Y)^\tau$ of all such consequence chains is denoted by $\mathbf{Y}(\mathbb{R} \times Y)$.

Let $\Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big)$ denote the set of all simple roulette lotteries over chains — that is, all finitely supported probability measures on $\mathbf{Y}(\mathbb{R} \times Y)$.

A *consequence history* $h \in H$ is a consequence chain which we will treat as having been realized before the initial node of the decision tree being considered.

2.2 The domain of timed consequence trees

The unrestricted domain $\mathcal{T}(\mathbb{R} \times Y)$ consists of all risky finite decision trees T whose set of nodes $N(T)$ can be partitioned into three classes:

$$N(T) = N^d(T) \cup N^c(T) \cup N^y(T),$$

where $N^d(T)$ denotes the set of *decision* nodes, $N^c(T)$ the set of *chance* nodes, and $N^y(T)$ the set of *timed consequence* nodes.

For each decision node $n \in N^d(T)$ there is a non-empty set $N^{+1}(n)$ of *immediate successor* nodes.

Given any finite set S, let $\Delta^0(S)$ denote the set of probability measures on S that have full support. For each chance node $n \in N^c(T)$ there is a specified roulette lottery $\pi(\cdot \mid n) \in \Delta^0\big(N^{+1}(n)\big)$ that selects the immediate successor of n.

Each timed consequence node $n \in N^y(T)$ is mapped to a timed consequence

$$\boldsymbol{y}(n) = (t(n), y(n)) \in \mathbb{R} \times Y.$$

Timed consequence nodes are partitioned into terminal consequence nodes

$$N^t(T) := \big\{n \in N^y(T) : N^{+1}(n) = \emptyset\big\}$$

and non-terminal (intermediate) consequence nodes $N^i(T) := N^y(T) \setminus N^t(T)$, where each $n \in N^i(T)$ has a unique successor. Thus, there is no branching at any consequence node, whether terminal or not.

The initial node of a decision tree T is denoted by $n_0(T)$. For any tree T and node $n \in N(T)$, the continuation subtree with initial node n is denoted by $T_{\geq n}$.

2.3 Path strategies and induced lotteries over consequence chains

Given any finite tree $T \in \mathcal{T}(\mathbb{R} \times Y)$, a *decision strategy* specifies, at each decision node $n \in N^d(T)$, a single successor n^{+1} in the set $N^{+1}(n)$ of all immediate successors of n. Any decision strategy generates an associated *path strategy* which, for each decision node n of T, specifies what unique lottery over entire paths through $T_{\geq n}$ from the initial node n to a random terminal node results from following the decision strategy at every decision node along a chosen path. Let $\mathbf{P}(T)$ denote the set of all path strategies in T.

Following a path strategy $\mathbf{p} \in \mathbf{P}(T)$ therefore generates a unique *random path* through the tree, and so a unique roulette lottery over timed consequence chains which is specified by the *path consequence mapping*

$$\mathbf{P}(T) \ni \mathbf{p} \mapsto \boldsymbol{\eta}(T, \mathbf{p}) \in \Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big) \tag{1}$$

Given any finite decision tree T in the unrestricted domain, and any decision strategy in T, in H&TM it is shown how to construct by backward recursion both the path strategy $\mathbf{p} \in \mathbf{P}(T)$ and the resulting lottery $\boldsymbol{\eta}(T, \mathbf{p})$ over consequence chains.

2.4 Feasible sets of lotteries over consequence chains

Given any node $n \in N(T)$ and any continuation subtree $T_{\geq n} \in \mathcal{T}(\mathbb{R} \times Y)$, let $F(T_{\geq n})$ denote the corresponding set of lotteries $\boldsymbol{\lambda} \in \Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big)$ over consequence chains $\mathbf{y}_{\geq n}$ that are feasible in the tree T because they result from a lottery $\boldsymbol{\eta}(T, \mathbf{p})$ induced by a path strategy $\mathbf{p} \in \mathbf{P}(T_{\geq n})$ in the continuation subtree $T_{\geq n}$.

In H&TM it is shown how in principle one can construct by backward recursion the entire family $\{F(T_{\geq n})\}_{n \in N(T)}$ of feasible sets that starts at each terminal node n^t of $T_{\geq n}$ and ends back at the initial node of T. When it reaches a chance node, the backward recursion involves a probability mixture of sets which is defined in the usual manner. Specifically, given any set of weights $(\alpha_i)_{i=1}^k$ such that $\alpha_i \geq 0$ for all i and $\sum_i \alpha_i = 1$, as well as corresponding sets $F_i \subseteq \Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big)$, we consider the corresponding convex combination

$$\sum_{i=1}^k \alpha_i F_i := \left\{ \sum_{i=1}^k \alpha_i \boldsymbol{\lambda}_i : \boldsymbol{\lambda}_i \in F_i \text{ for all } i \in \{1, \ldots, k\} \right\}.$$

2.5 Histories, behaviour rules and admissible consequences

Given any tree $T \in \mathcal{T}(\mathbb{R} \times Y)$ and any node $n \in N(T)$, there is a unique path of successive connected edges through the tree T that runs from its initial node $n_0(T)$ to n. This path passes through a unique finite sequence of timed consequence nodes, and so induces a unique consequence chain $\mathbf{y}_{\geq n_0(T),\, \leq n}$ that accrues along this unique path as one passes from $n_0(T)$ up to (but not including) node n. This allows one to associate with each node $n \in N(T)$ a uniquely determined history $h_{<n} = h_{<n_0(T)} \circ \mathbf{y}_{\geq n_0(T),\, \leq n} \in H$ consisting of the finite chain of timed consequences that results from prepending the history $h_{<n_0(T)}$ to the chain $\mathbf{y}_{\geq n_0(T),\, \leq n}$.

Let $\mathcal{T}^d(\mathbb{R} \times Y)$ denote the family of decision trees in the domain $\mathcal{T}(\mathbb{R} \times Y)$ whose initial node $n_0(T)$ is a decision node.

We need to describe behaviour at the initial node of each tree $T \in \mathcal{T}^d(\mathbb{R} \times Y)$ not only as a function of T, but also as a function of the previous history h before the initial node. This takes the form of a preceding consequence chain $h = \mathbf{y} \in H = \mathbf{Y}$.

Given any history $h \in H$, let $\mathcal{D}_h$ denote the family of all trees $T \in \mathcal{T}^d(\mathbb{R} \times Y)$ satisfying the restriction that the last timed consequence in the preceding consequence chain h occurs strictly before the first timed consequence in the tree T. Then the domain of the behaviour rule we are about to define is the set

$$\mathcal{D} := \{(h, T) \in H \times \mathcal{T}^d(\mathbb{R} \times Y) : T \in \mathcal{D}_h\} \tag{2}$$

After these preliminaries, we define a *behaviour rule* as a correspondence

$$\mathcal{D} \ni (h, T) \mapsto \beta_h(T) \in 2^{N^{+1}(n_0(T))} \setminus \{\emptyset\} \tag{3}$$

which selects a non-empty subset of the set $N^{+1}(n_0(T))$ of nodes resulting from moves that are immediately available at the initial node $n_0(T)$.

Note that any decision node n of T is the initial node of the continuation subtree $T_{\geq n} \in \mathcal{T}^d(\mathbb{R} \times Y)$, and that the pair $(h_{<n}, T_{\geq n})$ belongs to the domain $\mathcal{D}$. So, given the updated history $h_{<n}$, the same correspondence also specifies behaviour at n via its value $\beta_{h_{<n}}(T_{\geq n})$ at that initial node n of the tree $T_{\geq n}$.

Relative to a behaviour rule $\mathcal{D} \ni (h, T) \mapsto \beta_h(T)$, a path strategy $\mathbf{p} \in \mathbf{P}(T)$ is said to be *admissible* just in case, at every decision node n that is reached with strictly positive probability under $\mathbf{p}$, and for each possible history $h_{<n} \in H$, the immediate successor that is selected by $\mathbf{p}$ lies in the behaviour set $\beta_{h_{<n}}(T_{\geq n})$ prescribed by the behaviour rule β for the relevant continuation subtree $T_{\geq n}$.

Given any pair $(h, T) \in \mathcal{D}$, let $\mathbf{P}^{\beta}_h(T) \subseteq \mathbf{P}(T)$ denote the set of all path strategies in T that are admissible given the history h, and let

$$\Phi^{\beta}_h(T) := \left\{\boldsymbol{\eta}(T, \mathbf{p}) : \mathbf{p} \in \mathbf{P}^{\beta}_h(T)\right\} \;\subseteq\; F(T). \tag{4}$$

denote the corresponding set of *attainable consequences*.

2.6 Normal-form invariance, and prerationality

In H&TM and the previous papers cited there, a central principle is that, for each history $h \in H$ and each decision tree $T \in \mathcal{T}_h$, the consequences of behaviour in T should depend only on the relevant feasible set $F(T)$ of lotteries over timed consequence chains, and not on the particular dynamic structure of the tree T. This is formalised by the property of *consequentialist normal-form invariance*. Given any history $h \in H$, this requires that if the two trees $T, \widetilde{T} \in \mathcal{T}_h$ have the same feasible set $F_h(T) = F_h(\widetilde{T})$ contingent on the history h, then the behaviour rule should also induce the same admissible consequence set $\Phi^{\beta}_h(T) = \Phi^{\beta}_h(\widetilde{T})$.

If the behaviour rule β satisfies consequential normal form invariance, then given any history $h \in H$, the consequences of behaviour can be described equivalently by a conditional choice function C_h^β defined on finite feasible sets of lotteries over timed consequence chains. That is, consequential normal form invariance requires that, for each $h \in H$, there exists a non-empty-valued consequence choice function

$$\mathcal{F}(\Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big) \ni F \mapsto C_h^\beta(F) \in 2^F \setminus \{\emptyset\}$$

defined on the domain $\mathcal{F}\left(\Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big)\right)$ of finite subsets of $\Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big)$ such that, for every tree $T \in \mathcal{D}_h$, one has

$$\Phi_h^\beta(T) = C_h^\beta\big(F(T)\big) \subseteq \Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big).$$

We can then go on to define the *base preference relation* $\succsim_h^\beta$ that the choice function C_h^β induces on $\Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big)$ through the equivalence

$$\boldsymbol{\lambda} \succsim_h^\beta \boldsymbol{\lambda}' \iff \boldsymbol{\lambda} \in C_h^\beta\big(\{\boldsymbol{\lambda}, \boldsymbol{\lambda}'\}\big).$$

Following H&TM we say that, for each history $h \in H$, a base preference relation $\succsim_h$ on the domain $\Delta\big(\mathbf{Y}(\mathbb{R} \times Y)\big)$ is *prerational* just in case there exists a consequentialist normal-form invariant behaviour rule β_h on the unrestricted domain $\mathcal{T}(\mathbb{R} \times Y)$ that generates it, in the sense that $\succsim_h = \succsim_h^\beta$, where $\succsim_h^\beta$ is the induced base preference defined from the corresponding choice function C_h^β.

At this stage we can build on the results set out in Hammond (1998, 1998, 2022) and apply the main representation result of our previous paper H&TM. This states that, for each $h \in H$, the conditional base preference relation $\succsim_h$ on the unrestricted domain $\Delta(\mathbf{Y}(\mathbb{R} \times Y))$ of lotteries over chains of timed consequences is both prerational and mixture continuous if and only if it can be represented by the expected value of each conditional Bernoulli utility function $\mathbf{Y} \ni \mathbf{y} \mapsto u_h(\mathbf{y}) \in \mathbb{R}$ in a unique cardinal equivalence class.

3 Menu Consequences in Deterministic Decision Trees

This and the next three sections set out a series of examples that illustrate how allowing *menu consequences* substantially enlarges the kinds of behaviour that can be accommodated within the consequentialist framework. Theorem 8.1 implies that, whenever the underlying base preference on $\Delta(\mathbf{S}_\infty)$ is prerational and mixture continuous, behaviour in any tree $T \in \mathcal{T}(C)$ can be described as maximising the expected utility of the induced lottery over enriched consequence chains. The examples we provide exploit this observation in order to *pre-rationalise* a range of behaviour patterns that, when analysed within a narrower consequence domain, appear prima facie non-consequentialist.

The examples in this first of four consecutive sections involve deterministic decision trees in which, as illustrated in panel (b) of Figure 1, the typical path through the tree consists of four successive nodes:

1. first, there is an initial decision node;
2. second, there is a menu consequence node which has a menu attached;
3. next, there is a second decision node at which the feasible set corresponds to the menu attached to the previous node;
4. the fourth and last node is a terminal consequence node.

3.1 Preference for Flexibility

As far as we are aware, the two papers by Koopmans (1950, 1964) provide the first formal treatment of the idea of *preference for flexibility.*[1] Later Kreps (1979) provided axioms that characterize a preference to keep *options open.*

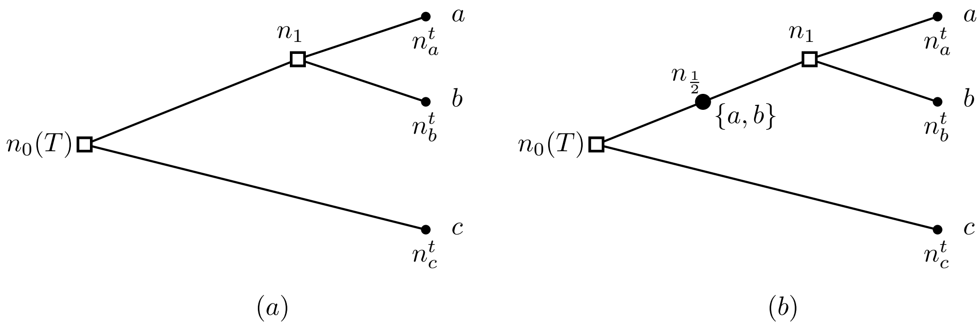

Figure 1: Rationalizing a preference for flexibility with menu consequences.

Figure 1 provides a simple illustration of how menu consequences can capture an intrinsic value of flexibility within a consequentialist framework. Starting from panel (a), the new panel (b) is obtained by inserting, just before the second decision node n_1, an added consequence node $n_{\frac{1}{2}}$ to which a menu consequence $M = \{a, b\}$ is attached. The resulting feasible consequence chains from the upper branch in panel (b) are (M, a) and (M, b), whereas in panel (a) they are simply a and b.

In this setting, it is natural to say that the decision maker exhibits a *preference for flexibility* whenever the experience of facing the menu $M = \{a, b\}$ strictly improves each continuation, in the sense that

$$(\{a, b\}, a) \succ a \quad \text{and} \quad (\{a, b\}, b) \succ b.$$

[1]So far we have been unable to find even a working paper version of the earlier work. We cite the abstract published in *Econometrica*.

This captures the idea that having the option set or menu $\{a, b\}$ available is itself valuable, even holding fixed the eventual continuation.

An alternative formulation closer in spirit to Kreps (1979) takes preferences over menus as primitive and imposes monotonicity w.r.t. set inclusion. In our framework this means that, for any two sets A and B whose members are consequence chains, one has

$$A \subseteq B, k \in A \implies (B, k) \succsim (A, k),$$

Thus, enlarging the menu weakly improves the resulting consequence chain, conditional on choosing the same continuation consequence chain k.

Finally, starting with the decision tree illustrated in panel (a) of Figure 1, introducing in panel (b) the menu consequence $\{a, b\}$ reverses the original preferences $c \succ a$ and $c \succ b$ in panel (a) in the sense that, in panel (b), one has

$$(\{a, b\}, a) \succ c \succ a \quad \text{and} \quad (\{a, b\}, b) \succ c \succ b.$$

If these strict preferences hold, we may say that there is a *strong preference for flexibility*. In panel (a) the commitment option c is preferred to either a or b, yet in panel (b) the upper branch is preferred because the flexibility premium attached to the menu consequence $\{a, b\}$ outweighs any advantage of commitment.

3.2 Choosing the Second Largest Piece of Cake

An even simpler example concerns deterministic choice under social norms. Sen (1997) discusses situations in which the decision maker avoids the objectively best available option, such as the most comfortable chair or the second largest piece of cake, in order not to appear impolite or even greedy. Using the same tree as in Figure 1, we interpret consequence a as the largest and consequence b as the second largest piece of cake.

In the absence of menu consequences, it is natural to have $a \succ b$, reflecting a taste for more cake. With menu consequences, however, the social significance of the choice can enter through the realised option set. If being offered the menu $M = \{a, b\}$ makes it salient that choosing a would be perceived as greedy, then the relevant ranking is no longer between a and b alone, but between the two chains (M, a) and (M, b). In particular, politeness can be represented by the preferences $(\{a, b\}, b) \succ (\{a, b\}, a)$ and, more generally, by the requirement that for any menu M containing a and b, one should have $(M, b) \succ (M, a)$ even though $a \succ b$ when the same outcomes are evaluated outside the context of the menu M.

3.3 The Potential Addict's Preference for Inflexibility

Hammond (1976) considers a "potential addict's decision tree", as illustrated in Fig. 2.

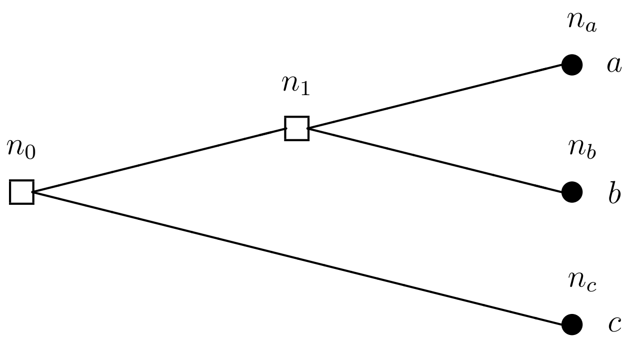

Figure 2: The potential addict's decision tree

The potential addict faces three possible consequences a, b, c, where a is addiction, b is bliss, and c is caution. One interprets b to mean that the agent benefits from taking the addictive substance for a short time, while it remains pleasurable, but then gives it up before harmful addiction sets in.

Initially at decision node n_0, the potential addict's preferences are $b \succ_0 c \succ_0 a$. But after taking the substance and reaching node n_1, addiction has set in so and the preference between a and b becomes $a \succ_1 b$. After including the menu consequence $\{a, b\}$ at the intermediate node $n_{1/2}$, the decision tree becomes that shown in Fig. 3. Say that there is:

- a *preference for flexibility* just in case both $(\{a, b\}, a) \succ a$ and $(\{a, b\}, b) \succ b$;
- a *preference for inflexibility* just in case both $a \succ (\{a, b\}, a)$ and $b \succ (\{a, b\}, b)$.

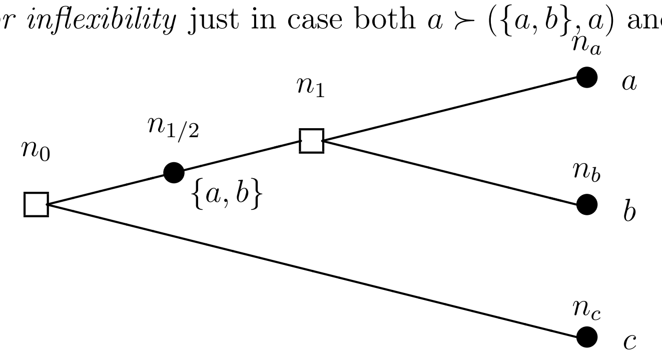

Figure 3: A Modified Tree T^M Admitting the Menu Consequence $\{a, b\}$

In the case of the potential addict, the preference $b \succ (\{a, b\}, b)$ may reflect a desire to avoid temptation.

3.4 Valuing an Individual's Freedom to Choose and Individual Rights

Sen (1970b) introduced individual rights into social choice theory. He modelled them as binary menus over which an individual's strict preference should be respected by a social choice rule.

The book by Rowley and Peacock (1975) formalized the idea that an individual's welfare could be enhanced by policies that are "liberal" in the sense of offering increased "freedom to choose", to quote a phrase popularized by Milton Friedman. See Sugden (1986, 2004, 2016) in particular for later work in this tradition.

A related way to model this freedom to choose was pioneered by Gibbard (1972), who uses game forms which relate to the "effectivity functions" introduced by Moulin and Peleg (1982). These functions specify sets of consequences that are menu consequences, in effect, because one individual has the power to determine which consequence in that set emerges as the social choice. The paper by Gaertner, Pattanaik, and Suzumura (1992) compared this new approach to Sen's. The two papers Hammond (1995, 1997) extended the comparisons but also advocated, in effect, that individual rights should be consequentialized.

3.5 Agency, Capabilities and Functionings

Sen (1985a) initiated a philosophical theory of individual agency. In Sen (1985b, 1993, 1999, 2002) as well as Nussbaum and Sen (1993), Drèze and Sen (2002), and Nussbaum (2011), this theory was applied in order to develop related concepts of capabilities and functionings. These were used as indicators in empirical studies intended to measure how much agency the individuals in different societies might have. Surveys appear in Sabina (2008) as well as in Welzel and Inglehart (2010).

Robeyns and Byskov (2020), in their encyclopedia entry, write that

> "Capabilities are the doings and beings that people can achieve if they so choose — their opportunity to do or be such things as being well-nourished, getting married, being educated, and travelling; functionings are capabilities that have been realized."

In this, they implicitly suggest that this wide variety of concepts related to agency could be regarded as different aspects of rich menu consequences.

3.6 Arrow's Rationalization for Violating Contraction Consistency

The following is based on an example which Kenneth Arrow reportedly offered orally in order to suggest that there could be circumstances which justify violating Chernoff's (1954) axiom, which Sen (1970a) described as "condition α" and others call "contraction consistencey".[2]

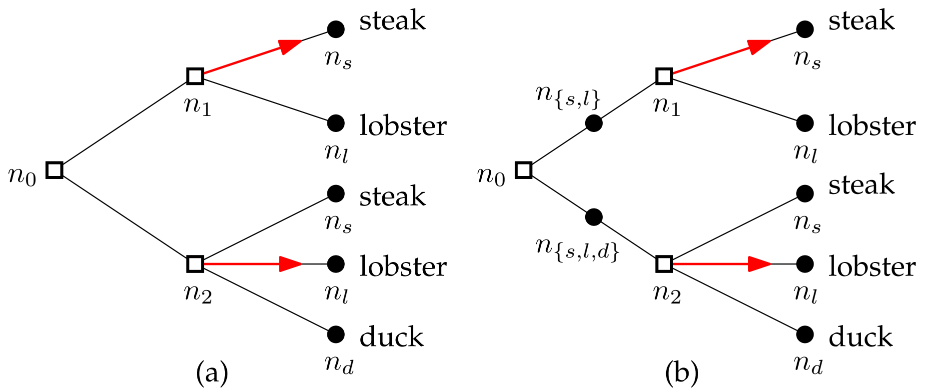

Figure 4: Arrow's example of violating contraction consistency

The example goes as follows. Suppose our agent goes to a fancy restaurant which advertises that *steak*, *lobster* and *duck* are on the menu. These are the three consequences shown at node n_2 in panel (a) of Figure 4. Here *steak* should be understood as a tasty but rather standard dish that one might find in many restaurants. On the other hand, *lobster* is a delicate dish that has to be not only fresh but also cooked with a high degree of skill.

Faced with the opportunity of having any one of the three delicacies in the feasible set $\{lobster, \; steak, \; duck\}$, our DM plans to choose *lobster*, as indicated by a red arrow emerging from node n_2 in panel (a) of Figure 4. After a mercifully brief delay, however, the waiter comes back with some bad news: they have run out of *duck*. With this additional information, our DM winds up ordering *steak* rather than *lobster* from the reduced menu $\{steak, \; lobster\}$, as shown by a red arrow emerging from node n_1 in panel (a) of Figure 4.

The two choices at nodes n_1 and n_2 in panel (a) of Figure 4 violate contraction consistency. Arrow's rationalization goes as follows: the fact that the restaurant has run out of *duck* signals that it is not that fancy. After all, what fancy restaurant runs out of *duck*? Possibly also signalling that they don't have their act together. Given that *lobster* is such a delicate dish, our agent doesn't want to risk it, given the extra information received, and goes for the safer option *steak*.

[2]We thank Herakles Polemarchakis for drawing our attention to this example.

Panel (b) of Figure 4 results from introducing the two consequence nodes $n_{\{s,l\}}$ and $n_{\{s,l,d\}}$ into the decision tree, with respective menu consequences

$$F(T_{\geq n_{\{s,l\}}}) = F(T_{\geq n_1}) = \{s, \ell\} \quad \text{and} \quad F(T_{\geq n_{\{s,l,d\}}}) = F(T_{\geq n_2}) = \{s, \ell, d\} \tag{5}$$

With the two menu consequences we have added, the relevant feasible set of consequence chains consists of the four possibilities

$$(\{s, \ell\}, s), \ (\{s, \ell\}, \ell), \ \{s, \ell, d\}, s), \ (\{s, \ell, d\}, \ell)$$

The extra items make it perfectly possible to have strict preferences that satisfy

$$(\{s, \ell\}, s) \succ (\{s, \ell\}, \ell) \quad \text{and yet} \quad (\{s, \ell, d\}, s) \prec (\{s, \ell, d\}, \ell)$$

These preferences, of course, rationalize choosing *lobster* in case *duck* is also available,

3.7 Conclusions

Menu consequences are quite flexible, and allow one to pre-rationalize many well known violations of contraction consistency, such as the *asymmetric dominance effect* first introduced by Huber et al. (1982). This effect occurs in an environment in which consumption consequences are described by their attributes. In the simplest case of two attributes, each consumption consequence $\mathbf{a}$ can be described by a pair $(a_1, a_2) \in \mathbb{R}_{\geq 0}$. Suppose that, when faced with the menu $\{\mathbf{a}, \mathbf{b}\}$, in which $a_1 > b_1$ and $a_2 < b_2$, the choice is alternative $\mathbf{a}$. The asymmetric dominance effect happens when a third alternative $\mathbf{c}$ is added to the menu, which is dominated by $\mathbf{b}$ (and not $\mathbf{a}$), that is, $c_1 < b_1$ and $c_2 < b_2$, and the resulting choice from $\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$ is alternative $\mathbf{b}$. It can readily be seen that a decision tree like that in Figure 4 can pre-rationalise this behavioural pattern with menu consequences, since it is pefcertly possible to have $(\{\mathbf{a}, \mathbf{b}\}, \mathbf{a}) \succ (\{\mathbf{a}, \mathbf{b}\}, \mathbf{b})$ and $(\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}, \mathbf{b}) \succ (\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}, \mathbf{a})$.

In conclusion, we mention the handbook chapter by Barberà, Bossert, and Pattanaik (2004) which surveys the literature on ranking "sets of objects", or menus.

4 Menu Consequences in Risky Decision Trees

4.1 Jacques Drèze's 1971 "Bergen Paradox"

The following example was the subject of considerable discussion, some of which the elder author of this paper had the privilege of witnessing. during the International Economic Association's (IEA) Workshop in Economic Theory held in Bergen, Norway, during the summer of 1971. A version was published on pages 14–15 of Drèze (1987), having appeared earlier in the volume devoted to part of the workshop that Drèze (1974) edited.

The example involves lotteries with random outcomes that belong to the consequence domain $Y = \{y_+, y_-, y_B, y_M\}$, where:

- y_+ is winning \$1000 (in 1971 dollars);
- y_- is losing \$1000 (in 1971 dollars);
- y_B is a ticket to a performance of Beethoven's 9th Symphony, whose last movement is a setting of Schiller's "Ode to Joy", which has become the EU's "national" anthem;
- y_M is a ticket to a performance of Mozart's Requiem.

For each $y \in Y$, let δ_y denote the degenerate roulette lottery in which consequence y occurs with probability one. Suppose $p, q \in (0,1)$, and consider the following four lotteries in $\Delta(Y)$:

$$\begin{aligned} \lambda_{+,B} &= p\,\delta_{y_+} + (1-p)\,\delta_{y_B} & \qquad \lambda_{+,M} &= p\,\delta_{y_+} + (1-p)\,\delta_{y_M} \\ \lambda_{-,B} &= q\,\delta_{y_-} + (1-q)\,\delta_{y_B} & \qquad \lambda_{-,M} &= q\,\delta_{y_-} + (1-q)\,\delta_{y_M} \end{aligned}$$

Consider also the following strict "mood matching" preferences over these lotteries:

- $\lambda_{+,B} \prec \lambda_{+,M}$ because winning a concert ticket is disappointing relative to winning \$1000, so $y_M \succ y_B$ if you are sad because you know you have not won \$1000;
- $\lambda_{-,B} \succ \lambda_{-,M}$ because winning a concert ticket might cause elation as opposed to losing \$ 1000, so $y_B \succ y_M$ if you know you have not lost \$1000.

Alternatively, there are also the following strict "mood damping" preferences over these lotteries:

- $\lambda_{+,B} \succ \lambda_{+,M}$ because winning a concert ticket is disappointing relative to winning \$1000, so $y_B \succ y_M$ if you need cheering up after learning you have not won \$1000;
- $\lambda_{-,B} \prec \lambda_{-,M}$ because winning a concert ticket might cause elation as opposed to losing \$ 1000, so $y_M \succ y_B$ because hearing a requiem is bearable if you know you have not lost \$1000.

The independence axiom, however, and so prerationality, requires:

$$\lambda_{+,B} \precsim \lambda_{+,M} \iff \delta_{y_B} \precsim \delta_{y_M} \iff \lambda_{-,B} \precsim \lambda_{-,M}$$

So both mood matching and mood damping preferences violate independence.

Suppose, however, that we treat each of the four lotteries $\lambda_{+,B}, \lambda_{+,M}, \lambda_{-,B}, \lambda_{-,M}$ as a possible menu consequence with the property that experiencing the lottery affects the decision maker's mood after the lottery has been resolved. Then feasible consequence chains combine:

1. one of the four lotteries $\lambda_{+,B}, \lambda_{+,M}, \lambda_{-,B}, \lambda_{-,M}$ as a menu consequence;
2. followed by whichever of the four consequences y_+, y_-, y_B, y_M is the random outcome of the lottery that is included in the menu consequence.

Introducing menu consequences in this way allows the suggested preferences between y_B and y_M that violate independence to be replaced by "mood matching" preferences

$$(\lambda_{+,B},\, y_B) \prec (\lambda_{+,M},\, y_M) \quad \text{and} \quad (\lambda_{-,B},\, y_B) \succ (\lambda_{-,M},\, y_M)$$

These allow the preference between y_B and y_M to depend on whether the alternative was winning or losing \$1000.

4.2 Lottery-Dependent Expected Utility

In the next example we include one menu consequence in an otherwise completely standard decision tree containing a single roulette lottery. We show that introducing this menu consequence allows us to pre-rationalise *prima facie* non-expected utility preferences. Consider a decision tree like that in Figure 5 which starts with the usual initial decision node n_0, followed by a chosen consequence node $n_{\{\ell\}}$ with a menu consequence consisting of the feasible set $F(T_{\geq n_\ell}) = \{\ell\}$ of consequences at the subtree following the node $n_{\{\ell\}}$. The third node on a typical path through the decision tree is a chance node at which the roulette lottery $\ell \in \Delta(Y)$ is resolved. The fourth node is a terminal node, one for each consequence in the set $\operatorname{supp} \ell = \{y, y'\}$.

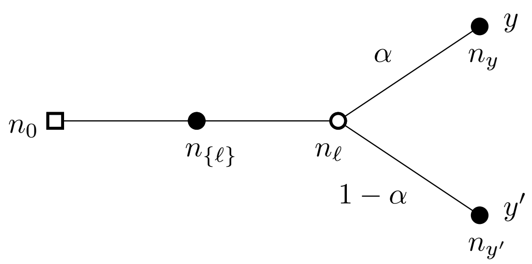

Figure 5: The roulette lottery $\ell = \alpha y + (1-\alpha) y'$ as a menu consequence at node $n_{\{\ell\}}$.

The two consequence chains corresponding to the two possible paths $(n_0, n_{\{\ell\}}, n_\ell, n_y)$ and $(n_0, n_{\{\ell\}}, n_\ell, n_{y'})$ shown in Fig. 5 are $(\{\ell\}, y)$ and $(\{\ell\}, y')$ respectively. So, for the unique feasible path strategy $\mathbf{p}$ at the only decision node n_0 in T, the corresponding lottery over consequence chains is $\boldsymbol{\lambda}(\mathbf{p}) = \alpha\,(\{\ell\}, y) + (1-\alpha)\,(\{\ell\}, y')$. Any resulting expected utility representation takes the form

$$U(\boldsymbol{\lambda}(\mathbf{p})) = \alpha\, u(\{\ell\}, y) + (1-\alpha)\, u(\{\ell\}, y')$$

In the subdomain of decision trees considered here, a strategy path first selects a roulette lottery $\ell \in \Delta(Y)$ and then chance determines its realised outcome $y \in Y$. The only menu consequence that appears on the chosen path through the tree is the singleton menu $\{\ell\}$, which therefore records which roulette lottery has been chosen. Hence every realised enriched consequence chain has the form $(\{\ell\}, y)$. Then, after identifying each singleton menu $\{\ell\}$ with the unique element $\ell \in \Delta(Y)$ it contains, the relevant consequence domain of pairs $(\{\ell\}, y)$ is naturally identified with $\Delta(Y) \times Y$.

Accordingly, prerationality and mixture continuity together imply that behaviour maximises the expected value of a Bernoulli index $\Delta(Y) \times Y \ni (\lambda, y) \mapsto u(\lambda, y) \in \mathbb{R}$ which is unique up to positive affine transformations.

Suppose $\mathbf{p}$ is a feasible path strategy which selects the roulette lottery $\lambda_{\mathbf{p}} \in \Delta(Y)$. Then the induced roulette lottery over enriched chains assigns probability $\lambda_{\mathbf{p}}(y)$ to each stream $(\{\lambda_{\mathbf{p}}\}, y)$. The corresponding expected utility evaluation is therefore

$$U\big(\boldsymbol{\lambda}(\mathbf{p})\big) = \sum_{y \in \operatorname{supp}(\lambda_{\mathbf{p}})} \lambda_{\mathbf{p}}(y)\, u(\lambda_{\mathbf{p}}, y). \tag{6}$$

It is important to emphasise that in a purely static choice domain a lottery-dependent representation such as (6) carries no testable implications apart from ordinality. If an analyst observes only choice behaviour over $\Delta(Y)$, then any preference representable by a real-valued utility function on $\Delta(Y)$ can be written in the form (6) by absorbing all dependence on λ into $u(\lambda, \cdot)$. What gives (6) substantive meaning in our setting is that it is obtained from behaviour on a richer dynamic domain, in which the realised stream explicitly records both the chosen roulette lottery and its subsequent realisation.

An amusing special case is when there exists a function $v : \Delta(Y) \to \mathbb{R}$ such that $u(\lambda, y) = v(\lambda)$ for all $(\lambda, y) \in \Delta(Y) \times Y$, so that the Bernoulli index is independent of the realised outcome y. Then, for any path strategy selecting λ, the induced expected utility collapses to

$$\sum_{y \in \operatorname{supp}(\lambda)} \lambda(y)\, u(\lambda, y) = \sum_{y \in \operatorname{supp}(\lambda)} \lambda(y)\, v(\lambda) = v(\lambda)$$

In this case, maximising expected utility over enriched chains is equivalent to maximising the (generally non-linear) functional v directly over $\Delta(Y)$. Thus, even though the representation takes an expected utility form on $\Delta(Y) \times Y$, the induced maximand over roulette lotteries in $\Delta(Y)$ need not be expected utility.

In this case, maximisating expected utility over the enriched domain $\Delta(\Delta(Y) \times Y)$

reduces to maximising v over $\Delta(Y)$. In other words, once menu consequences are admitted, a criterion that is *non-expected utility* when viewed as a preference over roulette lotteries $\lambda \in \Delta(Y)$ becomes a *special case* of expected utility when viewed instead as a preference over enriched consequence chains $(\lambda, y) \in \Delta(Y) \times Y$. The two statements concern different domains, but the message is that enriching the consequence description can consequentialise behaviour that would otherwise be classified as non-expected utility.

4.3 Preferences for Resolving a Lottery Sooner Versus Later

In this section we show that once menu consequences are admitted, our expected utility criterion on the domain of lotteries over consequence chains can pre-rationalise arbitrary preferences regarding the time at which a given roulette lottery is resolved. Kreps and Porteus (1978) introduced a canonical formalisation of such preferences over timing by working with a recursively specified class of dynamic choice problems, and then imposing axioms directly on the resulting recursive preference structure. A central motivation for their approach is their critique of what they call the *payoff vector approach*, requiring the decision maker's preferences to be represented by a Bernoulli utility function defined on chains of dated payoffs $(y_0, \ldots, y_T)$, with each strategy evaluated only through the probability distribution it induces over such payoff chains. In this framework, any two path strategies that induce the same distribution over payoff chains are necessarily indifferent, so there is no way to express a strict preference between any pair of lotteries that differ only in the *time at which uncertainty is resolved.* Indeed, Kreps and Porteus (1978) show that their generalisation collapses to the payoff-vector approach if and only if the decision maker is indifferent to when lotteries are resolved.

The point of the present example to be made here is that our framework is also a payoff-vector approach, but with payoff vectors given by *enriched consequence chains*. When menu consequences are allowed, these chains can record, as a payoff-relevant component, the continuation feasible set faced along the path. As a result, early and late resolution can induce different lotteries in the relevant domain even when they induce the same distribution over chains of elementary consequences. So an expected utility function on the domain of lotteries over extended consequence chains is compatible with strict preferences over the time at which lotteries are resolved.

Panel *(a)* of Figure 6 reproduces the canonical tree that Kreps and Porteus consider. At the initial decision node n_0, the decision maker chooses between the two roulette lotteries ℓ^E and ℓ^L. If the decision maker at node n_0 goes up to node n_{ℓ^E}, the resulting lottery ℓ^E resolves uncertainty early, in the sense that with probability α the agent

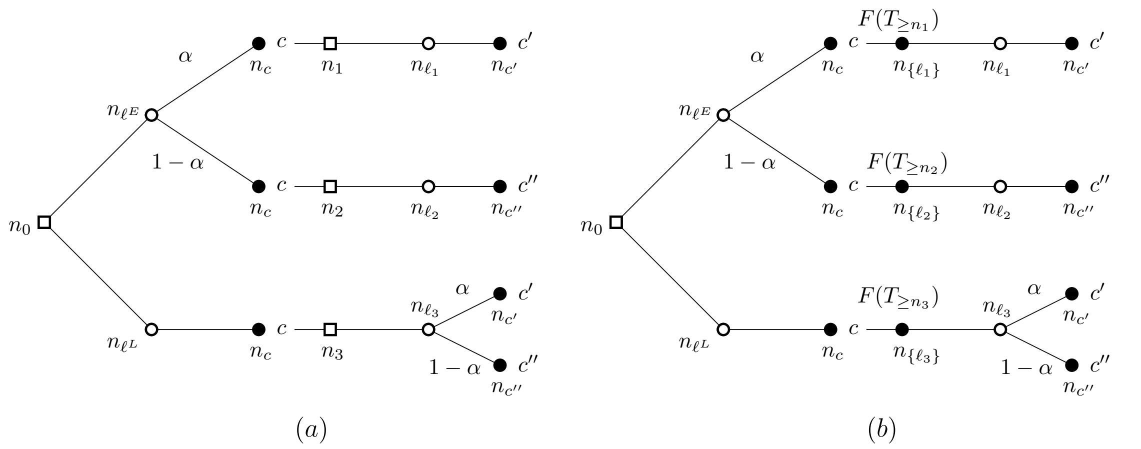

Figure 6: Panel *(a)* depicts the early versus late resolution tree of Kreps and Porteus (1978). Panel *(b)* depicts the corresponding representation with menu consequences.

learns in period 1 that period 2 consumption will be c', while with probability $1-\alpha$ she learns in period 1 that period 2 consumption will be c''. The alternative lottery ℓ^L resolves uncertainty late: in period 1 the agent only learns the distribution over $\{c', c''\}$, whereas the lottery is only realized in period 2. Noting that both lotteries ℓ^E and ℓ^L induce a probability distribution over pairs of the form *(consumption today, continuation lottery tomorrow)*, we may write

$$\ell^E = \alpha\,(c, \{\delta_{c'}\}) + (1-\alpha)\,(c, \{\delta_{c''}\}), \qquad \ell^L = (c, \{\alpha\delta_{c'} + (1-\alpha)\delta_{c''}\}).$$

In particular, both lotteries induce the same distribution $\alpha(c, c') + (1-\alpha)(c, c'')$ over deterministic consumption chains, since they differ only in when information is resolved.

Panel *(b)* shows how the distinction between early and later resolution of the lottery can be represented within our framework by inserting menu consequences. The singleton continuation "menus" n_1, n_2, n_3 in panel *(a)* get replaced by menu consequence nodes $n_{\{\ell_1\}}, n_{\{\ell_2\}}, n_{\{\ell_3\}}$, each carrying as a timed menu consequence the relevant continuation feasible set $F(T_{\geq n_i})$. In the tree shown in panel *(b)* of Fig. 6, these sets are precisely

$$F(T_{\geq n_1}) = \{\delta_{c'}\}, \qquad F(T_{\geq n_2}) = \{\delta_{c''}\}, \qquad F(T_{\geq n_3}) = \{\alpha\delta_{c'} + (1-\alpha)\delta_{c''}\}.$$

Let $\mathbf{p}^E$ and $\mathbf{p}^L$ denote the two path strategies corresponding resepectively to choosing ℓ^E and ℓ^L at n_0. Then the induced lotteries over enriched consequence chains are

$$\begin{aligned} &\boldsymbol{\eta}(T, \mathbf{p}^E) = \alpha\,(c, \{\delta_{c'}\}, c') + (1-\alpha)\,(c, \{\delta_{c''}\}, c'') \\ \text{and} \quad &\boldsymbol{\eta}(T, \mathbf{p}^L) = \alpha\,(c, \{\alpha\delta_{c'} + (1-\alpha)\delta_{c''}\}, c') + (1-\alpha)\,(c, \{\alpha\delta_{c'} + (1-\alpha)\delta_{c''}\}, c''). \end{aligned}$$

These two lotteries coincide after projecting chains onto their consumption coordinates, but they differ in their menu component. Hence an expected utility preference on the relevant domain can rank early versus late resolution by assigning different utilities to the intermediate menu consequences, without requiring a recursive primitive domain. In this sense, admitting menu consequences allows timing attitudes to be consequentialised, so that they become representable within a payoff-vector approach once payoff vectors are taken to be enriched consequence chains rather than consumption chains alone.

4.4 Two-period Epstein–Zin Preferences

Epstein and Zin (1989) postulated preferences which are designed to separate attitudes towards risk from attitudes towards intertemporal substitution or fluctuation in consumption levels. In their familiar recursive formulation, this separation is achieved by evaluating continuation risk through a certainty equivalent, and then aggregating certainty equivalents across time using a distinct curvature parameter that represents fluctuation aversion. In a two-period setting, the same idea can be expressed without recursion, and then it can be embedded in our framework once menu consequences are admitted.

Consider a two-period risky decision tree in which:

1. first, the decision maker initially chooses a first-period consumption roulette lottery $\lambda_1 \in \Delta(C_1)$;
2. next, the lottery λ_1 is resolved to yield consumption $c_1 \in C_1$;
3. next, the decision maker chooses a second-period consumption roulette lottery $\lambda_2 \in \Delta(C_2)$;
4. finally, the lottery λ_2 is resolved to yield consumption $c_2 \in C_2$.

In order to allow fluctuation aversion to be distinguished from risk aversion, we enrich the tree by inserting, immediately before each of the two chance nodes at which a roulette lottery is resolved, a menu-consequence node that records the singleton menu $\{\lambda_t\}$ for $t = 1, 2$. As in Section 4.2, we identify each realised singleton menu $\{\lambda_t\}$ with its unique element λ_t. This allows us to regard the relavant consequence in each period $t \in \{1, 2\}$ as the pair $(\lambda_t, c_t) \in \Delta(C_t) \times C_t$. Thus the relevant consequence domain for each two-period risky decision tree is $(\Delta(C_1) \times C_1) \times (\Delta(C_2) \times C_2)$. In particular, each enriched consequence chain we have constructed records for each period both the chosen roulette lottery and its realised consumption level.

Let $\rho > 0$ with $\rho \neq 1$ be the consumer's constant relative rate of risk aversion. Given the two consumption lotteries (λ_1, λ_2) define their *certainty equivalent* consumption levels $c_t^E(\lambda_t)$ implicitly by

$$\frac{1}{1-\rho}\Big(c_t^E(\lambda_t)\Big)^{1-\rho} = \sum_{c\in\operatorname{supp}(\lambda_t)} \lambda_t(c)\,\frac{1}{1-\rho}\,c^{\,1-\rho}. \tag{7}$$

Equivalently, one has

$$c_t^E(\lambda_t) = \Bigg(\sum_{c\in\operatorname{supp}(\lambda_t)} \lambda_t(c)\,c^{\,1-\rho}\Bigg)^{\frac{1}{1-\rho}}$$

Next let $\phi > 0$ with $\phi \neq 1$ be the consumer's constant relative rate of fluctuation aversion, or constant elasticity of intertemporal substitution. Let $\beta \in (0,1)$ be the consumer's discount factor, and define the two-period Epstein–Zin evaluation of the pair (λ_1, λ_2) by

$$V(\lambda_1,\lambda_2) := \frac{1}{1-\phi}\Big[\big(c_1^E(\lambda_1)\big)^{1-\phi} + \beta\,\big(c_2^E(\lambda_2)\big)^{1-\phi}\Big]. \tag{8}$$

When $\rho = \phi$, Equation (8) reduces to the standard time separable expected-utility criterion

$$\frac{1}{1-\rho}\Big[\sum_{c_1}\lambda_1(c_1)c_1^{1-\rho} + \beta\sum_{c_2}\lambda_2(c_2)c_2^{1-\rho}\Big]$$

When $\rho \neq \phi$, however, the evaluation (8) is not expected utility over the joint distribution of (c_1, c_2) because the risk in each period is first replaced by a certainty equivalent using ρ, before being aggregated across time using ϕ.

In our framework, however, (8) can be represented by expected utility on the enriched stream domain. To see this, define a Bernoulli utlity function on enriched chains by

$$u\big((\lambda_1,c_1),(\lambda_2,c_2)\big) := \frac{1}{1-\phi}\Big[\big(c_1^E(\lambda_1)\big)^{1-\phi} + \beta\,\big(c_2^E(\lambda_2)\big)^{1-\phi}\Big],$$

of the certainty equivalents in each of the two periods. This depends on the realised consumption chain only through the two menu consequences λ_1 and λ_2. Then, for any path strategy that selects (λ_1, λ_2), the induced lottery over enriched chains assigns probability $\lambda_1(c_1)\lambda_2(c_2)$ to the stream $((\lambda_1, c_1), (\lambda_2, c_2))$. So expected utility yields

$$\sum_{c_1,c_2}\lambda_1(c_1)\lambda_2(c_2)\,u\big((\lambda_1,c_1),(\lambda_2,c_2)\big) = V(\lambda_1,\lambda_2).$$

Thus two-period Epstein–Zin preferences can be pre-rationalised in our framework by treating the chosen roulette lotteries as menu consequences. From the perspective of the consumption outcomes alone, the resulting criterion is generally non-expected utility when $\rho \neq \phi$. Yet it is expected utility on $\Delta(\mathbf{S}_\infty)$, because the enriched consequence stream allows its menu components to record payoff-relevant features of the timing and structure of the decision problem.

4.5 Temptation and Self-Control

Gul and Pesendorfer (2001) model temptation by taking *menus of lotteries* as primitive objects of choice. Their key behavioural assumption, which they call *set betweenness*, requires that whenever a menu A is weakly preferred to a menu B, then the enlarged menu $A \cup B$ is ranked in between, in the sense that

$$A \succsim B \implies A \succsim A \cup B \succsim B,$$

This creates the logical possibility that the agent could be made worse off with extra unchosen feasible options because they create temptation. Gul and Pesendorfer's representation of the agent's preferences uses two indices: a normative utility u and a temptation utility v. Then the *self-control cost* of choosing $y \in A$ is measured by the $v-$opportunity cost $\max_{x \in A} v(x) - v(y)$ relative to the most tempting alternative in A.

Figure 7 shows how our framework can consequentialise the same idea. The agent first reaches a consequence node n_0 with a menu consequence in the form of the set $F(T_{\geq n_0}) = \{\ell, \ell'\} \subseteq \Delta(Y)$ of consequence lotteries that are feasible in the continuation subtree $T_{\geq n_0}$. At the immediately succeeding decision node n_1 the choice is between the two continuation paths leading to:

1. either the consequence node $n_{\{\ell\}}$, followed by the chance node n_ℓ at which the random outcome of the lottery $\ell = \alpha \delta_y + (1 - \alpha)\, \delta_{y'}$ is determined;

2. or the consequence node $n_{\{\ell'\}}$, followed by the chance node $n_{\ell'}$ at which the random outcome of the lottery $\ell' = \beta \delta_y + (1 - \beta)\, \delta_{y'}$ is determined.

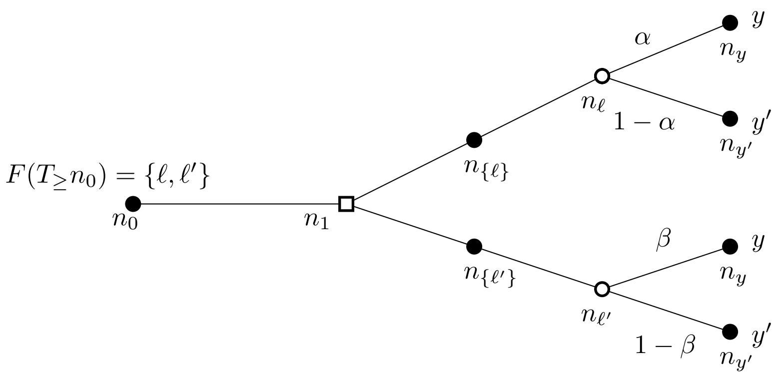

Figure 7: Introducing the menu $F(T_{\geq n_0}) = \{\ell, \ell'\}$ consequentializes temptation.

In terms of enriched consequence chains, choosing ℓ from the menu $\{\ell, \ell'\}$ and then realising $y \in Y$ yields the consequence chain $(\{\ell, \ell'\}, \ell, y)$. Hence the path strategy $\mathbf{p}_\ell$

that chooses ℓ at n_1 induces the lottery

$$\boldsymbol{\eta}(T, \mathbf{p}_\ell) = \alpha\delta_y + (1-\alpha)\,\delta_{y'}$$

and similarly for $\mathbf{p}_{\ell'}$. Then prerationality and mixture continuity together imply that there exists a Bernoulli index $(\{\ell,\ell'\},\ell,y) \mapsto u(\{\ell,\ell'\},\ell,y) \in \mathbb{R}$ on realised chains whose expectation is a von Neumann–Morgenstern utility function which behaviour maximizes.

To recover Gul and Pesendorfer's self-control criterion, fix two Bernoulli utility functions $\hat{u}, \hat{v} : Y \to \mathbb{R}$ on the domain Y of simple consequences. Then define the Bernoulli index on the space of realised chains by

$$u(\{\ell,\ell'\},\ell,y) := \hat{u}(y) - \Big(\max_{k\in\{\ell,\ell'\}} \mathbb{E}_k[\hat{v}] - \mathbb{E}_\ell[\hat{v}]\Big) \tag{9}$$

and analogously for $u(\{\ell,\ell'\},\ell',y)$. Taking expectations of (9) w.r.t. the probabilities specified by the lottery ℓ gives

$$\mathbb{E}_\ell[u(\{\ell,\ell'\},\ell,\tilde{y})] = \mathbb{E}_\ell[\hat{u}] - \Big[\max_{k\in\{\ell,\ell'\}} \mathbb{E}_k[\hat{v}] - \mathbb{E}_\ell[\hat{v}]\Big] = \mathbb{E}_\ell[\hat{u}+\hat{v}] - \max_{k\in\{\ell,\ell'\}} \mathbb{E}_k[\hat{v}]$$

It follows that maximising expected utility over the set of feasible path strategies at n_1 is equivalent to

$$\max_{\ell^*\in\{\ell,\ell'\}} \Big\{\mathbb{E}_{\ell^*}[\hat{u}+\hat{v}]\Big\} - \max_{k\in\{\ell,\ell'\}} \mathbb{E}_k[\hat{v}],$$

This is exactly the Gul–Pesendorfer representation in which:

1. $\hat{u}$ plays the role of the commitment ranking;

2. $\hat{v}$ plays the role of the temptation ranking;

3. the term $\max_k \mathbb{E}_k[\hat{v}] - \mathbb{E}_{\ell^*}[\hat{v}]$ is the self-control cost of choosing ℓ^* in the presence of the most tempting alternative.

Finally, Gul and Pesendorfer also analyse an *extended* preference on pairs (A, x), where A is the menu chosen in period 1 and $x \in A$ is the option chosen in period 2. In Figure 7, the corresponding object is naturally identified with one of the two interim segments $(\{\ell,\ell'\},\ell)$ and $(\{\ell,\ell'\},\ell')$ that are parts of the corresponding consequence chains that have each been enriched with the appropriate menu consequence. Hence, by including the menu consequence at n_0, our framework can accommodate not only the ex ante ranking of menus, but also the within-menu temptation comparisons that, in the analysis by Gul and Pesendorfer, drive second-period choice.

4.6 History-Dependent Self-Control

Troccoli Moretti (2024) studies a dynamic model of choice under risk in which, the more disappointed the decision maker has been in the past, the stronger is the temptation she can face in the future. The *disappointment history-dependent self-control* (or D-HDSC) model that the paper introduces is set in the choice domain introduced by Kreps and Porteus (1978). As argued in the paper, in the absence of menu consequences, this implies that behaviour violates consequentialism because conditioning second-period preferences on past disappointments makes the agent care about counterfactual, unrealised branches of the tree. Yet, as we shall see in this section, introducing menu consequences allows D-HDSC preferences to be pre-rationalized.

To illustrate how this model works, consider the decision tree T depicted in Figure 8. It consists of an initial (degenerate) decision node n_0, followed by a menu consequence node n_ℓ at which the consequence is the menu $F(T_{\geq n_\ell}) = \{\ell\}$. Next comes a chance node n_1 with a single roulette lottery ℓ which delivers either high or low first-period consumption c^h or c^ℓ as a timed consequence in the first period. Next, after each of these two consequence nodes comes a second menu consequence node $n_{\{s,ns\}}$ with the appropriate menu consequence $F(T_{\geq n_1})$ or $F(T_{\geq n_2})$. There follows a second decision node n_1 or n_2, which offers the DM the choice between either smoking and going to a terminal node labelled n_s or not smoking and going to a terminal node labelled n_{ns}.

Observe that the tree T has three menu consequence nodes. Of these, the first is the menu consequence $n_{\{\ell\}} = F(T_{\geq} n_\ell)$ which implies that, in principle, when evaluating a particular branch of the tree, the agent may care about the structure of the continuation subtree $T_{\geq n_{\{\ell\}}}$. In particular, it is this menu consequence that allows the agent to care about histories of disappointment. Second, there are two subsequent menu consequence nodes labelled $n_{\{s,ns\}}$ with respective menu consequences $F(T_{\geq n_j})$ for $j \in \{1, 2\}$. Their presence allows the agent to care about menu effects, which in the present case manifest themselves by creating self-control problems.

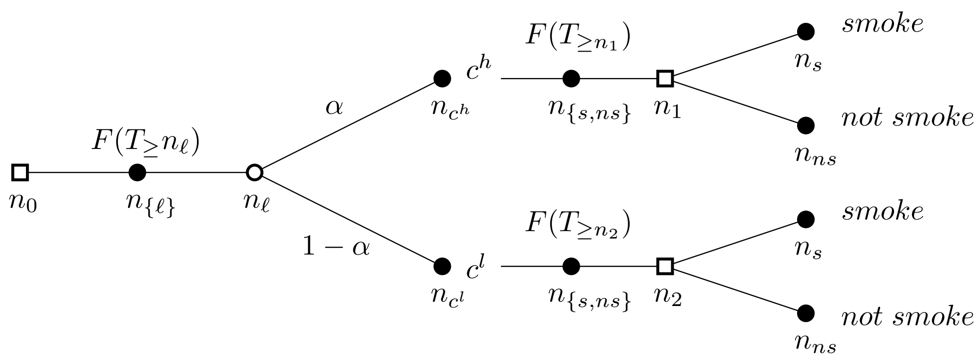

Figure 8: The decision tree T motivating D-HDSC preferences.

With the structure of the tree T shown in Fig. 8, a typical consequence chain takes the form $(\{\ell\}, c^k, \{s, ns\}, a)$ with $k \in \{h, l\}$ followed by an action $a \in \{s, ns\}$. Combining the non-trivial choices at the two nodes n_1 and n_2 in tree T gives four possible decision strategies. Observe that both decision nodes n_1 and n_2 offer the same choice between smoking and not smoking. The possibility of having experienced consequence lotteries, however, allows entire consumption histories to matter, and hence offers scope for behaviour at nodes n_1 and n_2 to differ.

Also, introducing menu consequences allows for the additional dependence on histories that contain unrealised branches of the tree. If disappointments are relevant, then not only does the history given by the branch connecting node n_0 to n_2 matter, but so also does the counterfactual edge of the tree that links node n_ℓ to n_{c^h}. Observe that adding menu consequences drastically expands the type of history dependence that is allowed. In fact, the type of history dependence considered in D-HDSC preferences even allows behaviour to differ at two decision nodes that offer the same choice and arise from the same history of past consumption. This is intuitive since facing a low consumption stream with certainty induces no disappointment, while getting a low consumption history as a result of a lottery does disappoint because the agent knows that things could have been better.

5 Menu Consequences with Uncertainty

5.1 Roulette versus Horse Lotteries

In their paper offering a definition of subjective probability, Anscombe and Aumann (1963) draw an important distinction between:

1. roulette lotteries with specified objective or hypothetical probabilities;

2. horse lotteries without specified objective or hypothetical probabilities.

In particular, if probabilities are defined for a horse lottery, they must be subjective. This distinction between roulette and horse Lotteries mirrors how most economists have come to understand the distinction between risk and uncertainty.

This Section is the only part of the paper which pays any attention to horse lotteries, even though they play a key role in the earlier work on consequentialism and prerationality that appears in, for example, Hammond (1998, 2022). This is because horse lotteries add little to our discussion of menu consequences. A significant exception, however, arises in the work by Ellsberg (1961) who, even before Anscombe and Aumann (1963) had appeared, used examples to show the importance in decision theory of the distinctions between not only roulette and horse lotteries, between risk and uncertainty, and between objective and subective probability. There is also an associated distinction between chance and event nodes in decision trees.

5.2 The Ellsberg Paradox

Ellsberg's (1961, p. 653) main example, which we are about to recapitulate, violates Anscombe and Aumann (1963) extension of Savage's sure-thing principle rather than, as Ellsberg claimed, the original principle itself.

Lottery	R	B	Y	$w(q)$
a	1	0	0	1/3
b	0	1	0	q
c	1	0	1	$1-q$
d	0	1	1	2/3

Table 1: Ellberg's Urn and the Probabilities $w(q)$ of Winning

The example we shall discuss involves an urn which is known to contain 90 balls, of which exactly 30 are red (R), whereas the remaining 60 are either black (B) or yellow (Y) without the exact numbers of each being known. One ball is drawn at random; all 90 balls are equally likely to be drawn. There are four lotteries $\ell \in \{a, b, c, d\}$ set out in the rows of Table 1 with consequences 0, which represents "no prize", or 1, which represents a fixed "prize".[3] As a function of q, the uncertain proportion of black balls, the last column of Table 1 reports for each of the four lotteries the conditional probability $w(q)$ of winning the prize.

To describe the options facing the agent, for each pair $F \subset \{a, b, c, d\}$ we consider the finite decision tree T^F in which:

1. first, the agent faces the decision node $\{n_0\}$ at which the decision is between the two immediately succeeding nodes n_ℓ^F ($\ell \in F$);
2. second, for each lottery $\ell \in F$, node n_ℓ^F is an event node at which a horse lottery determines which of the set $\{n_{\ell,q}^F\}_{q \in Q}$ of 61 subsequent nodes is reached, where:
   (a) $Q := (1/90)\,(\mathbb{Z} \cap [0, 60]) \subset [0, \frac{2}{3}]$ is the set of integer multiples of 1/90 for an integer in the interval $[0, 60]$;
   (b) each $q \in Q$ specifies what proportion of the 90 balls are black, implying that a proportion $\frac{2}{3} - q$ are yellow;

[3] Ellsberg (1961) specified that it was \$100, presumably at 1961 prices in the U.S. Also, he labelled the four urns using upper case Roman numerals rather than letters.

3. third, each $n^F_{\ell,q}$ is a chance node where a roulette lottery $\lambda^F_{\ell,q}$ determines which of the three possible random colours R, B, Y occurs, with specified probabilities $\frac{1}{3}$, q, and $\frac{2}{3} - q$.

Ellsberg reported that experimental subjects typically reported preferences satisfying $a \succ b$ and $d \succ c$. These violate Anscombe and Aumann's extension of Savage's sure-thing principle because, regardless of whether the "sure-thing" is that Y occurs or alternatively that it does not, the conditional horse lotteries given the pair $\{R, B\}$ are identical for a and c, as well as for b and d. So one should have $a \succ b \Longleftrightarrow c \succ d$. which is violated by the preferences that are typically reported.

Following Gärdenfors and Sahlin (1982), amongst others, an explanation that is often suggested is that a and d both reduce to roulette lotteries offering known and certain probabilities of winning the lottery, whereas b and c offer uncertain probabilities. In order to rationalize these preferences, we introduce into each decision tree T^F, and immediately prior to each succeeding chance node $n^F_{\ell,q}$, an additional menu consequence node $\hat{n}^F_{\ell,q}$ where the menu consequence is the singleton set $\{\lambda^F_{\ell,q}\}$.

Now, just as the Epstein–Zin preferences with menu consequences that were discussed in Section 4.4 allow different degrees of risk and fluctuation aversion, here menu consequences will allow different degrees of risk and uncertainty aversion. Indeed, we suppose that the agent is risk neutral but uncertainty averse. Specifically, we suppose that there exists a strictly convex cost function $[0, 1] \ni w \mapsto C(w)$ of the winning probability with the property that the agent's Bernoulli utility function of the lottery prize $p \in \{0, 1\}$ and the winning probability $w \in [0, 1]$ takes the form

$$\{0, 1\} \times [0, 1] \ni (p, w) \mapsto p - C(w) \in \mathbb{R}$$

Suppose that agent's subjective beliefs concerning the proportion q of black balls in the urn, and so of the associated probability $w(q)$ of winning any specified bet, take the form of the subjective probability density function $Q \ni q \mapsto \mathbb{P}_q \in [0, 1]$ satisfying $\sum_{q \in Q} \mathbb{P}_q = 1$. Under the assumptions stated above, the value of the agent's subjectively expected utility, which depends on the subjective expectations regarding the outcome $(q, w(q))$, is

$$U(\mathbb{P}) := \mathbb{E}_{\mathbb{P}} \left[w(q) - C\left(w(q)\right) \right]$$

In the case of the four bets set out in Table 1, with respective winning probabilities set out in the last column, the subjectively expected von Neumann–Morgenstern utilities, in an obvious notation, are

$$U_a = \tfrac{1}{3} - C(\tfrac{1}{3}), \quad U_b = \mathbb{E}_{\mathbb{P}}[q - C(q)], \quad U_c = \mathbb{E}_{\mathbb{P}}[1 - q - C(1 - q)], \quad U_d = \tfrac{2}{3} - C(\tfrac{2}{3}) \tag{10}$$

We make one final assumption, which is that $\mathbb{E}_{\mathbb{P}} q = \frac{1}{3}$. This is true if, for instance, the density function $Q \ni q \mapsto \mathbb{P}_q \in [0,1]$ is *symmetric* in the sense that $\mathbb{P}_q(q) = \mathbb{P}_q(\frac{2}{3} - q)$ for all $q \in Q$. Under this symmetry assumption, in the degenerate case when $\mathbb{P}(\frac{1}{3}) = 1$, and so there is no uncertainty about the probability w of winning, then evidently (10) implies that $U_a = U_b$ and $U_c = U_d$. Otherwise, applying Jensen's theorem for strictly convex functions to the expected utilities in (10) implies that

$$\begin{aligned} & U_b = \mathbb{E}_{\mathbb{P}}[q - C(q)] < \tfrac{1}{3} - C\left(\mathbb{E}_{\mathbb{P}} q\right) = \tfrac{1}{3} - C(\tfrac{1}{3}) = U_a \\ \text{and} \quad & U_c = \mathbb{E}_{\mathbb{P}}[1 - q - C(1-q)] < \tfrac{2}{3} - C\left(\mathbb{E}_{\mathbb{P}}(1-q)\right) = \tfrac{2}{3} - C(\tfrac{2}{3}) = U_d \end{aligned}$$

These inequalities, of course, correspond to the reported preferences $a \succ b$ and $d \succ c$.

5.3 The Precautionary Principle

Keynes (1930, 1936) apparently introduced the idea that one motive for saving might be precaution in the face of risk, uncertainty, or even ignorance regarding how much wealth should be saved in order to help provide for the future. The idea that precaution can be explained by a preference for a menu consequence that keeps options open for later is implicit in the work by Arrow and Fisher (1974), which considers the imperfectly forseeable option value of preserving an environmental resource. The case they make for precaution is related to, but seems more general than, the explanation in Hammond (2026) which concerns the subjective evaluations that a bounded rational agent might want to attach to the terminal nodes of a truncated decision tree.

6 Menu Consequences in Social Choice

Ever since the work of Bergson (1938), social choice theory has usually presumed that social welfare is a Paretian function of individual utility levels. So typically the main effect of allowing social choice to depend on menu consequences is due to the effect on individual well-being. Nevertheless, there are significant effects which arise when, for instance, the menu consequences are inherently social, as they are, for example, when they involve concepts like equality and social justice.

6.1 Machina's Mom

Machina (1989, p. 1643) provides an example of a Mom with an indivisible treat to allocate to one of her two small children A or B. Consider a decision tree where the initial decision node offers a choice between the three immediately succeeding nodes:

1. the terminal consequence node n_a, whose consequence is the degenerate lottery δ_A where A gets the treat with probability 1, while B bursts into tears;

2. the terminal consequence node n_b, whose consequence is the degenerate lottery δ_B where B gets the treat with probability 1, while A bursts into tears;

3. the intermediate consequence node n_{mix}, where there is a menu consequence given by the mixed lottery $\frac{1}{2}\delta_A + \frac{1}{2}\delta_B$, whose two respective successors are copies n'_a of n_a and n'_b of n_b.

Without the menu consequence, the independence axiom of EU theory allows Machina's Mom to have only one of the following three preference relations over the three lotteries in the consequence domain $\{\delta_A, \delta_B, \frac{1}{2}\delta_A + \frac{1}{2}\delta_B\}$:

$$\text{(i) } \delta_A \succ \tfrac{1}{2}\delta_A + \tfrac{1}{2}\delta_B \succ \delta_B; \quad \text{(ii) } \delta_B \succ \tfrac{1}{2}\delta_A + \tfrac{1}{2}\delta_B \succ \delta_A; \quad \text{(iii) } \delta_A \sim \tfrac{1}{2}\delta_A + \tfrac{1}{2}\delta_B \sim \delta_B.$$

So the fair lottery $\frac{1}{2}\delta_A + \frac{1}{2}\delta_B$ cannot be uniquely best, even though intuitively one might feel that it should be.

With a menu consequence, however, the relevant domain of consequence chains becomes $\{\delta_A, \delta_B, (\frac{1}{2}\delta_A + \frac{1}{2}\delta_B, \delta_A), (\frac{1}{2}\delta_A + \frac{1}{2}\delta_B, \delta_B)\}$. Then Machina's Mom can strictly prefer either of the two consequence chains which start with the lottery $\frac{1}{2}\delta_A + \frac{1}{2}\delta_B$ to both those consequence chains which do not.

6.2 Diamond's Critique of Harsanyi's Version of the Original Position

In a widely cited paper, Diamond (1967) offered a powerful critique of Harsanyi's (1953, 1955) justification of utilitarianism. The rest of this subsection presents a somewhat simplified but, at least in our view, no less persuasive version of Diamond's example. It argues that the force of Diamond's critique can be considerably reduced by adding a menu consequence which helps society distinguish between egalitarian and inegalitarian social states, and so to express a preference for egalitarianism.

Indeed, like Diamond, consider a society consisting of two individuals labelled A and B. Suppose that the set of possible social states consists of the four points labelled (y^A, y^B) that belong to the Cartesian product $Y^A \times Y^B$ where $Y^A = Y^B = Y = \{0, 1\}$. Consider the following two risky social states:

1. the *egalitarian* social state $\frac{1}{2}\delta_{(1^A,1^B)} + \frac{1}{2}\delta_{(0^A,0^B)}$ in which the two individuals' personal consequences y^A and y^B are perfectly correlated;

2. the *inegalitarian* social state $\frac{1}{2}\delta_{(1^A,0^B)} + \frac{1}{2}\delta_{(0^A,1^B)}$ in which the two individuals' personal consequences y^A and y^B are perfectly anti-correlated.

Now, assuming symmetry, Harsanyi's average measure of social welfare can be expressed as the average $W = \frac{1}{2}\mathbb{E}[u(y^A)] + \frac{1}{2}\mathbb{E}[u(y^B)]$ of the two individuals' expected values of a common utility function $Y \ni y \in u(y) \in \mathbb{R}$. So $W = \frac{1}{2}u(1) + \frac{1}{2}u(0)$ for both the egalitarian and the inegalitarian social states. Then Harsanyi's utilitarianism, according to this argument, regards both the egalitarian and inegalitarian social states as indifferent. Yet a common opinion is that, other things being equal, individuals have better lives if they belong to a more equal society.

With menu consequences, however, we can reconcile utilitarianism with egalitarianism in an obvious way. All we need to do is to consequentialize interpersonal equality. Some ways of doing this are described in the next subsection.

6.3 Rawlsian Maximin

Consider a society with a finite set of individuals $i \in I$, and for each $i \in I$ a common personal consequence domain Y_i. Consider any social decision tree with a single initial decision node n_0 and, for some finite set $F \subseteq \prod_{i\in I} Y_i$ of personal consequence profiles, an associated family $N^{+1}(n_0) = \{n_{y^I} \mid y^I \in F\}$ of terminal consequence nodes that immediately succeed n where, for each $y^I \in F$, the consequence of reaching each n_{y^I} is y^I.

Without any menu consequence, we follow Vickrey (1945) and Harsanyi (1953, 1955) in considering an *original position.* This takes the form of an extended decision tree for an *impartial benefactor* in which, for each interpersonal profile $y^I \in F$ of personal consequences, the terminal consequence node n_{y^I} of the original tree is replaced by a new chance node $\tilde{n}_{y^I}$ at which the consequence is the symmetric *impersonal lottery* $(1/\#I)\sum_{i\in I} \delta_{\tilde{n}_{y_i}}$ of experiencing each person i's randomized consequence y_i after emerging from behind a different version of what Rawls (1971) calls "the veil of ignorance". The appropriate social objective that is implied is the utilitarian sum $\sum_{i\in I} u(y_i)$.

Alternatively, suppose the impartial benefactor faces a decision tree with a menu consequence based on the Rawlsian original position. Specifically, suppose we can precede each chance node $\tilde{n}_{y^I}$ by the combination of:

1. a horse lottery whose outcome in each uncertain state of the world $i \in I$ is the relevant personal consequence $y_i \in Y_i$;

2. a preceding menu consequence which is this horse lottery.

In the decision tree with menu consequences, a more egalitarian social objective could be the modified utilitarian sum $\sum_{i\in I} v(\min_{i\in I}\{u(y_i)\}, y_i)$ that includes a Rawlsian evaluation $\min_{i\in I}\{u(y_i)\}$ of the horse lottery menu consequence in the impartial benefactor's assessment of each individual i's personal consequence y_i.

6.4 Prioritarianism

The term "prioritarianism" may well be due to Temkin (2003), though he gives most of the credit to the noted philosopher Derek Parfit. A definition due to Rabinowicz (2002) states that:

> "Prioritarianism applies some increasing concave transformation to a measure of utility that a utilitarian might want to use."

In its extreme form, following the results in Hammond (1975) applied to utilities rather than to incomes, it becomes either the maximin criterion discussed in Section 6.3 or, if discontinuities are allowed, the leximin criterion.

An alternative approach is to specify an additive social welfare function that is more egalitarian than the utilitarian sum $\sum_{i \in I} u(y_i)$ but less egalitarian than the "Rawlsian" function $\sum_{i \in I} v(\min_{i \in I}\{u(y_i)\}, y_i)$ suggested in Section 6.3. The idea is to replace the extremely inequality averse function $\min_{i \in I}\{u(y_i)\}$ of the menu consequence y^I in the original position by an alternative less inequality averse symmetric measure $y^{\mathrm{EDE}}(y^I)$ of equally distributed equivalent income, as proposed by Atkinson (1970). Indeed, given a strictly increasing and strictly concave utility function $\mathbb{R}_+ \ni y \mapsto u(y) \in \mathbb{R}$ of income, the equally distributed equivalent income is defined implicitly by the equation $u(y^{\mathrm{EDE}}) = (1/\#I) \sum_{i \in I} u(y_i)$. This of course is the analogue for the theory of income inequality of a certainty equivalent income in the theory of risk-taking. The final result is the "equality moderated" utilitarian social welfare function $\sum_{i \in I} v(y^{\mathrm{EDE}}(y^I), y_i)$.

7 Menus as Consequences

7.1 Complications When Menus are Consequences

Our previous paper H&TM showed that the earlier results regarding prerational behaviour in finite decision trees can be extended to allow finite chains of timed consequences. Given any non-empty domain Y of basic consequences, then as discussed in Section 2.1, the extension set out in H&TM requires the domain of possible lottery consequences to be extended from $\Delta(Y)$ to the space $\Delta(\mathbf{Y}(\mathbb{R} \times Y))$ of lotteries over chains $\mathbf{y} \in \mathbf{Y}(\mathbb{R} \times Y)$ of timed consequences. Here, the examples we have presented in the four Sections 3–6 illustrate how much might be gained by extending orthodox decision theory even further in order to allow the timed consequences of decisions to include menu consequences.

The goal of this Section is to investigate the complications that arise when trying to carry out this extension. A genuine conceptual and technical difficulty is that the domain of menu consequences must include finite menus of roulette lotteries over the entire domain of continuation chains. Yet each continuation chain may itself include menu consequences encountered later on the same path through the decision tree. This creates a circularity: the domain of feasible menu consequences depends on the domain of feasible consequence chains, which depends in turn on what sets of consequences may appear as timed consequences.

To resolve this circularity, we use an iterative procedure of *finite depth* in order to construct the outcome domain, taking as given an elementary consequence domain Y. The first step starts from the baseline space $\mathbf{S}_0 = \mathbf{Y}(\mathbb{R} \times Y)$ of chains of primitive timed consequences (t, y). Then, as explained in Section 7.2 and Appendix A, we iteratively build hierarchies of menu domains and enriched chain domains. Given the chains constructed up to depth k, one forms the finite menu domain $\mathbf{M}_k = \mathcal{F}(\Delta(\mathbf{S}_k))$. The next chain domain $\mathbf{S}_{k+1}$ consists of the self-verifying chains in $\mathbf{Y}(\mathbb{R} \times (Y \sqcup \mathbf{M}_k))$: whenever such a chain contains a menu consequence, the suffix that follows it must be a possible realisation of one of the lotteries in that menu. The last step constructs the limiting envelopes

$$\mathbf{S}_\infty := \bigcup_{k \geq 0} \mathbf{S}_k, \qquad \mathbf{M}_\infty := \bigcup_{k \geq 0} \mathbf{M}_k.$$

As shown in Appendix A, this construction gives $\mathbf{M}_\infty = \mathcal{F}(\Delta(\mathbf{S}_\infty))$, and every menu-consistent finite decision tree generates feasible lotteries over the self-verifying chain domain $\mathbf{S}_\infty$.

A *menu consequence* makes this object itself payoff-relevant by allowing the decision tree to record, at some date, the menu of continuation prospects that is available from that point onwards. Thus, in addition to primitive timed consequences (t, y) with $y \in Y$, we allow timed consequences of the form (t, F), where F is a finite menu of lotteries over appropriately enriched consequence chains.

Consider first a deterministic tree with decision nodes and timed consequence nodes only. Fix a node n. Any strategy path for the continuation from n uniquely determines a complete timed consequence chain from n to a terminal node. It is therefore natural to regard the continuation feasible set at n as the set $F(T_{\geq n})$ of all consequence chains that can be generated in the continuation subtree that starts from n. When the tree includes chance nodes, this set becomes a finite menu of roulette lotteries over consequence chains. In the presence of history-dependent preferences, these lotteries are the appropriate objects of choice because different continuation chains may be ranked differently even when they share the same terminal primitive consequence.

Now suppose the tree contains a menu consequence node m. The menu consequence attached to m is intended to be the feasible set of consequences at the successor of m, interpreted as the menu of continuation lotteries that the agent faces after observing m. This is straightforward when menus are *simple* in the sense that no path through the tree can include more than one menu consequence node. In that case, each menu element can be identified with an ordinary continuation chain, or with a lottery over such chains when chance is present.

The complication arises when some paths can contain more than one menu consequence node. By design, the consequence attached to a menu node is a *menu of continuation prospects*, namely the feasible set of consequences from that point onwards, which takes the form of a finite menu of roulette lotteries over complete continuation chains. If the tree can include another menu node later on the same path, then some of those continuation chains no longer consist only of primitive timed consequences. Rather, they are chains that include, at the relevant dates, the menu objects encountered later on the path. Hence the domain must be closed under the operation "menus are sets of lotteries over chains" while simultaneously allowing "chains to contain menus as timed consequences". Put differently, to define a menu consequence we must already know what the feasible continuation menus are. But to define feasible continuation menus, we must already know what the admissible chain objects are, and these chain objects may themselves contain menus. This is the basic recursive dependence that a well-defined domain construction has to disentangle.

7.2 The Domain of Finite Decision Trees with Menu Consequences

We now explain how to incorporate menu consequences into the timed-consequence framework. The key modelling decision is that a *menu consequence node* is an intermediate timed consequence node. Thus it carries a time label but, as at any intermediate consequence node, there is no branching at that node. The unique successor of a menu consequence node determines what the menu label must represent.

7.2.1 What a menu consequence records

Fix an intermediate consequence node $n \in N^i(T)$ with unique successor $n^{+1} \in N^{+1}(n)$. The intended interpretation is that a menu consequence at n records, at time $t(n)$, the continuation opportunities that are available from the node that immediately succeeds n.

If n^{+1} is a decision node, then these opportunities are precisely the feasible set of continuation consequences in the continuation subtree $T_{\geq n^{+1}}$, namely the set $F(T_{\geq n^{+1}})$ of roulette lotteries over complete continuation chains.

If n^{+1} is a chance node with roulette probability kernel $\pi(\cdot \mid n^{+1})$, then the continuation opportunities depend on which random successor node $n' \in N^{+1}(n^{+1})$ is realised. One can describe the resulting menu consequence as either a roulette lottery over the family $\{F(T_{\geq n'})\}_{n' \in N^{+1}(n^{+1})}$ of continuation menus, or equivalently, by using the reduction of compound lotteries that is already built into the recursive construction of feasible sets. In our formal construction we follow the latter convention, so that menu consequences remain *menus of roulette lotteries over chains.*

Finally, if n^{+1} is itself a timed consequence node, then the continuation opportunities that the menu consequence is meant to record are determined only after one moves further along the unique consequence-only segment, until the first subsequent chance, event, or decision node is reached. Operationally, there is no loss in collapsing any sequence of consecutive menu consequences into a single one.

The common point across these cases is that a menu consequence node does not introduce any further new feasible continuation prospects. Rather, it records as a timed consequence the continuation feasible set that is already determined by what follows. The menu-consistency condition imposed in Appendix A makes this interpretation formal: the menu consequence attached to an intermediate menu-consequence node must be the feasible set generated by its successor subtree.

7.2.2 Menu trees and why the construction iterates

The appropriate way to think about multiple menu consequences on the same path is that they induce a finite hierarchy of *future menus*. Starting from any menu consequence node, consider what other menu consequence nodes may be met later along a path, and which further menu consequence nodes may be met later still. Since the underlying decision tree is finite, this hierarchy has finite depth: along any realised path, only finitely many menu consequence nodes can be encountered, and each earlier menu can only refer to menus that lie deeper in this hierarchy.

This hierarchical structure explains why the outcomes cannot be defined in a single step. The menu recorded at an earlier node must be a menu of roulette lotteries over *complete* continuations. But a complete continuation may include, at later dates, the menus that will be faced at later menu consequence nodes. Hence, to specify what can be an element of an earlier menu, one must already have specified what a later continuation chain is allowed to contain. The natural resolution is to construct a sequence of consequence domains indexed by depth. One begins with chains of depth 0 that contain only primitive timed consequences, and forms menus of depth 0 as finite sets of roulette lotteries over these chains. Next, one allows menus of depth 0 to appear as timed consequences inside chains of depth 1, forms menus over lotteries of depth 1 chains, and iterates. Taking the union over all finite depths yields the finite-depth enriched chain domain used below.

7.2.3 The finite-depth self-verifying construction

Let $\mathcal{F}(X)$ denote the family of all finite non-empty subsets of a set X. Our iterative construction starts with the depth 0 domains

$$\mathbf{S}_0 := \mathbf{Y}(\mathbb{R} \times Y) \quad \text{and} \quad \mathbf{M}_0 := \mathcal{F}\big(\Delta(\mathbf{S}_0)\big) \tag{11}$$

of depth 0 consequence chains, and of finitely supported roulette lotteries over $\mathbf{S}_0$. Given any $k \in \{0\} \cup \mathbb{N}$ and the pair $(\mathbf{S}_k, \mathbf{M}_k)$, the next step of the inductive construction defines $\mathbf{S}_{k+1}$ as the set of self-verifying chains in

$$\mathbf{Y}\big(\mathbb{R} \times (Y \sqcup \mathbf{M}_k)\big).$$

That is, a chain belongs to $\mathbf{S}_{k+1}$ only if, whenever it contains a menu consequence, the suffix that follows that menu is a possible realisation of one of the lotteries in the menu. The associated menu domain is then [4]

$$\mathbf{M}_{k+1} := \mathcal{F}\big(\Delta(\mathbf{S}_{k+1})\big). \tag{12}$$

Finally, define the two limits

$$\mathbf{S}_\infty := \bigcup_{k=0}^{\infty} \mathbf{S}_k \quad \text{and} \quad \mathbf{M}_\infty := \bigcup_{k=0}^{\infty} \mathbf{M}_k. \tag{13}$$

Thus $\mathbf{S}_\infty$ contains the finite self-verifying chains that can be generated by recursively nesting menus to any finite depth. A chain that records a menu M and then continues with a suffix that is not encoded by M is not an element of $\mathbf{S}_\infty$. Since every decision tree we allow is finite, each tree has a finite bound on the number of menu consequence nodes that can be encountered along any path. Appendix A gives the formal details of this construction, proves that $\mathbf{M}_\infty = \mathcal{F}(\Delta(\mathbf{S}_\infty))$, and shows that every menu-consistent finite decision tree generates feasible lotteries over $\mathbf{S}_\infty$.

7.2.4 The unrestricted domain with menu consequences

We therefore take the syntactic consequence domain to be $\mathcal{C} := \mathbb{R} \times \big(Y \sqcup \mathbf{M}_\infty\big)$. Then we define the unrestricted domain of risky finite decision trees with timed menu consequences to be

$$\mathcal{T}(\mathcal{C}) = \mathcal{T}\big(\mathbb{R} \times (Y \sqcup \mathbf{M}_\infty)\big).$$

[4] For any two sets A, B, we denote by $A \sqcup B$ their *disjoint union*, namely $(A \times \{0\}) \cup (B \times \{1\})$. We use disjoint union to keep primitive consequences and menu objects as different sorts, so that a timed consequence (t, z) unambiguously carries either $z \in Y$ or $z \in \mathbf{M}_k$, and never conflates the two.

A tree $T \in \mathcal{T}(\mathcal{C})$ has decision nodes, chance nodes, and timed consequence nodes, as before. Each intermediate consequence node $n \in N^i(T)$ carries a timed label

$$\boldsymbol{y}(n) = (t(n), z(n)) \in \mathbb{R} \times \big(Y \sqcup \mathbf{M}_\infty\big),$$

where $z(n) \in Y$ corresponds to an ordinary primitive timed consequence, whereas $z(n) \in \mathbf{M}_\infty$ corresponds to a menu consequence label.

The unrestricted class $\mathcal{T}(\mathcal{C})$ is larger than the domain used for behaviour. The behaviourally relevant dynamic domain is the history-indexed class $\mathcal{D}^m$ constructed in Appendix A. In that domain, terminal timed consequence nodes carry primitive labels, and every intermediate menu consequence records the feasible set generated by its successor subtree.

7.2.5 Feasible sets on the enlarged domain

For the new enriched consequence domain $\mathcal{C} := \mathbb{R} \times \big(Y \sqcup \mathbf{M}_\infty\big)$, the definition of path strategies is unchanged. For an arbitrary syntactic tree $T \in \mathcal{T}(\mathcal{C})$, the usual backward recursion formally assigns to each path strategy $\mathbf{p} \in \mathbf{P}(T)$ a roulette lottery over well-typed enriched chains. For the menu-consistent trees in $\mathcal{D}^m$, Appendix A shows that the realised chains belong to $\mathbf{S}_\infty$. Accordingly, for every $(h, T) \in \mathcal{D}^m$, the path strategy consequence mapping takes the form

$$\mathbf{P}(T) \ni \mathbf{p} \mapsto \boldsymbol{\eta}(T, \mathbf{p}) \in \Delta\big(\mathbf{S}_\infty\big).$$

Thus the feasible set $F(T)$ of consequences in any $T \in \mathcal{T}_h^m$ is a subset of the richer space $\Delta(\mathbf{S}_\infty)$ of lotteries. More generally, for any node n and its accrued history $h_{<n}$, the continuation feasible set $F_{h_{<n}}(T_{\geq n})$ is the set of roulette lotteries over $\mathbf{S}_\infty$ induced by path strategies in the continuation subtree $T_{\geq n}$, evaluated from the starting history $h_{<n}$. Behaviour rules and admissibility are also defined exactly as before, as well as the set $\mathbf{P}^\beta(T) \subseteq \mathbf{P}(T)$ of β-admissible path strategies, and the corresponding attainable set $\Phi^\beta(T) \subseteq F(T)$ of consequence chains.

The substantive step in extending the framework is therefore the construction of the self-verifying chain domain $\mathbf{S}_\infty$, together with the menu-consistent dynamic domain $\mathcal{D}^m$. Once these have been constructed, the definitions and structural characterisations of $\boldsymbol{\eta}$, $F(\cdot)$, and $\Phi^\beta(\cdot)$ extend as in H&TM, with $\mathbf{Y}(\mathbb{R} \times Y)$ replaced by $\mathbf{S}_\infty$ and $\Delta(\mathbf{Y}(\mathbb{R} \times Y))$ replaced by $\Delta(\mathbf{S}_\infty)$. The representation theorem in the next section is stated on this lottery domain.

8 Bayesian Rationality with Menu Consequences

This section extends the representation theorems in H&TM to the enriched domain that allows *menu consequences*. The objective is to obtain an analogue of Theorem 5.1 in H&TM: we take an arbitrary base preference relation on the relevant lottery domain and show that it is prerational and mixture continuous if and only if it admits an expected utility representation over finite consequence chains. In our setting, the relevant lottery domain is $\Delta(\mathbf{S}_\infty)$, where $\mathbf{S}_\infty$ is the self-verifying consequence-chain domain constructed in Section 7.2 and Appendix A. Chains in $\mathbf{S}_\infty$ may contain both timed elementary consequences and timed menu consequences, with menu nesting allowed at any finite depth.

The main proof strategy mirrors that in H&TM. Consequentialist normal-form invariance allows one to collapse any decision tree into a consequentially normal-form equivalent tree in which the DM chooses among path strategies, and each path strategy induces a lottery over consequence chains. This reduction turns the dynamic problem into a static choice problem over a finite menu of lotteries over chains, which is exactly the domain on which the earlier results operate. For the menu-consistent dynamic trees constructed in Appendix A, each path strategy induces a lottery in $\Delta(\mathbf{S}_\infty)$. For the binary comparisons used to recover the base preference on $\Delta(\mathbf{S}_\infty)$, the relevant trees are normal-form comparison trees whose terminal consequences are already constructed chains in $\mathbf{S}_\infty$.

Allowing menu consequences creates a complication connected with preference revelation. If a menu consequence F occurs before a later decision node, then the relevant base preference relation at that decision node is $\succsim_h$, where F is part of the history h. Thus the resulting preferences are menu-dependent. As is familiar from static choice domains with menu-dependent preferences, observed choice need not reveal the ranking of unchosen alternatives. For example, if $F = \{a, b, c\}$ is realised and the behaviour rule always selects a from the subsequent choice among $\{a, b, c\}$, this behaviour need not reveal how (F, b) and (F, c) are ranked. The dynamic structure of decision trees can recover this comparison: one may keep F as part of the realised history and insert a later decision node at which the choice is between b and c.

There is a second, distinct issue. Let $M = \{a, b\}$ and $N = \{c, d\}$. The chains (M, a) and (N, c) are legitimate self-verifying consequence chains. Each can be generated by a menu-consistent dynamic tree. But a dynamic tree that tried to compare them by first recording M and then offering only a, and first recording N and then offering only c, would not be menu-consistent. The first branch would make M false, and the second would make N false. The comparison is obtained by a binary normal-form tree: a tree whose initial node is a decision node with two available actions, one inducing the consequence chain (M, a) and the other inducing the consequence chain (N, c). Once (M, a) and (N, c) belong to $\mathbf{S}_\infty$, the unrestricted-domain hypothesis allows such a tree to be constructed.

This is the interpretation required by normal-form invariance. A dynamic choice problem is reduced to a static choice among strategies, and each strategy has an induced consequence chain. Thus it is coherent to ask the decision maker to compare the strategy "choose the upper branch and then a", whose induced consequence chain is (M, a), with the strategy "choose the lower branch and then c", whose induced consequence chain is (N, c). In normal form, these chains are consequences. The binary tree just described contains no internal menu-consequence node, so menu consistency is vacuous.

Proposition 1 (Prerationality implies completeness)**.** *For any history $h \in H$, let $\succsim_h$ be a base preference relation on $\Delta(\mathbf{S}_\infty)$. If $\succsim_h$ is prerational, then $\succsim_h$ is complete on $\Delta(\mathbf{S}_\infty)$.*

The proof of Proposition 1, given in Appendix B, uses normal-form comparison trees over the already constructed chain domain $\mathbf{S}_\infty$. For any two lotteries $\boldsymbol{\lambda}, \boldsymbol{\mu} \in \Delta(\mathbf{S}_\infty)$, one considers a binary normal-form tree whose feasible set is $\{\boldsymbol{\lambda}, \boldsymbol{\mu}\}$. This tree presents lotteries over consequence chains as terminal consequence objects; it does not re-expand those chains into new dynamic menu-consequence trees.

The earlier example with $F = \{a, b, c\}$ shows how the dynamic structure recovers comparisons between unchosen alternatives inside a realised menu. Since F remains part of the history, a later decision node offering b and c induces the chains (F, b) and (F, c), and behaviour at that node reveals the missing binary comparison.

The example with $M = \{a, b\}$ and $N = \{c, d\}$ shows why the general completeness argument uses normal-form comparisons over already constructed consequence chains. Menu consistency constrains dynamic menu-consequence trees; it does not require normal-form comparison trees to re-expand their terminal chains into new dynamic menu trees. Once completeness is recovered on $\Delta(\mathbf{S}_\infty)$, the usual consequentialist arguments yield transitivity and independence on $\Delta(\mathbf{S}_\infty)$.

We therefore obtain the following analogue of Theorem 5.1 in H&TM for the enriched domain that allows menu consequences.

Theorem 8.1. *The base preference relation $\succsim$ on $\Delta(\mathbf{S}_\infty)$ is prerational and mixture continuous if and only if there exists a Bernoulli utility index $\mathbf{S}_\infty \ni \mathbf{y} \mapsto u(\mathbf{y}) \in \mathbb{R}$ such that the von Neumann–Morgenstern expected utility function $\Delta(\mathbf{S}_\infty) \ni \boldsymbol{\lambda} \mapsto U(\boldsymbol{\lambda}) = \sum_{\mathbf{y} \in \operatorname{supp}(\boldsymbol{\lambda})} \boldsymbol{\lambda}(\mathbf{y})\, u(\mathbf{y}) \in \mathbb{R}$ satisfies*

$$\boldsymbol{\lambda} \succsim \boldsymbol{\lambda}' \iff U(\boldsymbol{\lambda}) \geq U(\boldsymbol{\lambda}') \tag{14}$$

Moreover, the Bernoulli utility index u belongs to a unique cardinal equivalence class

$$\mathsf{A}_u := \Big\{ v \in \mathbb{R}^{\mathbf{S}_\infty} : \exists (a,b) \in \mathbb{R}_{>0} \times \mathbb{R} \text{ such that } u = av + b \Big\}.$$

Theorem 8.1 therefore implies that, whenever the base preference relation is prerational and mixture continuous, the decision maker chooses among feasible path strategies by maximising the expected utility of the induced lottery over enriched consequence chains, where those chains may include menu consequences. For every $(h, T) \in \mathcal{D}^m$, Appendix A shows that $F(T) \subseteq \Delta(\mathbf{S}_\infty)$. Hence the representation applies to all feasible path strategies in every menu-consistent dynamic tree. The binary comparisons used in the proof are normal-form comparisons over already constructed consequence chains; the evaluation of dynamic trees uses the feasible lotteries generated by menu-consistent trees.

9 Concluding Summary

By adapting earlier results from Hammond (1988, 1998), it was shown in Hammond (2022) that, for finite decision trees that have only terminal consequences, prerationality combined with mixture continuity of preferences implies that a decision maker's decisions should be Bayesian rational, meaning that they maximize the expected value of a Bernoulli utility function. The main result of H&TM is that the same conclusion holds for decision trees that admit chains of consequence nodes, of which at most one can be terminal. This paper has extended the scope of Bayesian rationality further to include decision trees in which any non-terminal consequence node n may have a menu consequence in the form of the continuation subtree whose initial node is n.

The examples presented in Sections 3–6 illustrate the very rich variety of possibilities that arise once decision trees are allowed to have consequence nodes to which menu consequences are attached. Indeed, introducing menu consequences can help explain patterns of behaviour that a decision theorist might otherwise regard as anomalous. Such consequences can make these anomalies seem plausible, even prerational. Rationality, however, goes beyond mere prerationality in requiring also that a decision-maker should act as if pursuing rational goals, whatever that may mean. Evidently this consideration limits, perhaps rather severely, what can rationally be counted as a relevant menu consequence.

Indeed, one criticism of consequentialization discussed by Brown (2011) is that it is in danger of becoming tautological. This, of course, is a devastating critique of a theory that is meant to be purely descriptive. But much of our work concerns prescriptive theory. Then, if a consequentialist theory has become tautological, this might even be seen as an advantage because it encourages a search beyond mere prerationality for more refined criteria that distinguish a good decision theory from a bad one.

Much work still remains to be done in exploring the range of new possibilities.

Appendix A The Domain of Finite Decision Trees with Menu Consequences

This appendix constructs the domain of finite risky decision trees whose intermediate timed consequences may be menu consequences. The construction has two parts. First, it defines the enriched stream and menu domains $(\mathbf{S}_\infty, \mathbf{M}_\infty)$. These domains solve the recursive dependence created by menu consequences: a menu consequence is a menu of lotteries over continuation streams, but those continuation streams may themselves contain later menu consequences. The stream domain $\mathbf{S}_\infty$ is restricted to self-verifying streams, so that whenever a stream contains a menu consequence, the suffix that follows it is a possible realisation of one of the lotteries in that menu. Second, the appendix defines histories as prefixes of self-verifying streams, and constructs the history-compatible and menu-consistent tree classes $\mathcal{T}_h$ and $\mathcal{T}_h^m$. These classes ensure that timed consequences occur after the relevant history, that histories generated along the tree remain admissible prefixes, and that each menu consequence records the feasible set generated by the continuation subtree that follows it.

A.1 Finite enriched streams and menus

Let Y be a non-empty set of elementary consequences. A timed consequence chain over $\mathbb{R} \times Y$ is a finite sequence

$$\mathbf{y} = \big((t_1, y_1), \ldots, (t_\tau, y_\tau)\big) \in (\mathbb{R} \times Y)^\tau$$

such that $\tau \in \mathbb{N}$ and $t_1 < \cdots < t_\tau$. We write

$$\mathbf{Y}(\mathbb{R} \times Y) := \bigcup_{\tau \in \mathbb{N}} \big\{ (t_j, y_j)_{j=1}^\tau \in (\mathbb{R} \times Y)^\tau : t_{j+1} > t_j \text{ for all } j = 1, \ldots, \tau - 1 \big\}$$

for the set of all finite timed consequence chains over $\mathbb{R} \times Y$.

For any set X, let $\Delta(X)$ denote the set of finitely supported roulette lotteries over X, and let $\mathcal{F}(X)$ denote the family of all finite non-empty subsets of X. For any two sets A and B, let

$$A \sqcup B := (A \times \{0\}) \cup (B \times \{1\})$$

denote their disjoint union. We write $\iota_A : A \to A \sqcup B$ and $\iota_B : B \to A \sqcup B$ for the canonical injections. The disjoint union keeps elementary consequences and menu consequences formally distinct: an element of Y and an element of a menu domain are never identified as the same consequence. For a finite stream

$$\mathbf{y} = \big((t_1, z_1), \ldots, (t_\tau, z_\tau)\big),$$

and any $j < \tau$, write

$$\mathbf{y}_{>j} := \big((t_{j+1}, z_{j+1}), \ldots, (t_\tau, z_\tau)\big)$$

for the suffix following the j-th timed consequence. Also, for any finite stream $\mathbf{s}$ and any $t \in \mathbb{R}$, write $\mathbf{s} > t$ if every date appearing in $\mathbf{s}$ is strictly greater than t.

We now construct the enriched stream and menu domains recursively. Set

$$\mathbf{S}_0 := \mathbf{Y}(\mathbb{R} \times Y), \qquad \mathbf{M}_0 := \mathcal{F}\big(\Delta(\mathbf{S}_0)\big).$$

Thus $\mathbf{S}_0$ is the baseline domain of finite timed consequence chains with no menu consequences, whereas $\mathbf{M}_0$ is the collection of finite menus of finitely supported lotteries over such chains. Given $\mathbf{M}_k$ where $k \in \{0\} \cup \mathbb{N}$, define $\mathbf{S}_{k+1}$ as the set of all finite timed streams

$$\mathbf{y} = \big((t_1, z_1), \ldots, (t_\tau, z_\tau)\big) \in \mathbf{Y}\big(\mathbb{R} \times (Y \sqcup \mathbf{M}_k)\big)$$

such that, whenever $z_j = M \in \mathbf{M}_k$, the following conditions hold:

$$j < \tau, \qquad \text{there exists } \boldsymbol{\lambda} \in M \text{ such that } \mathbf{y}_{>j} \in \operatorname{supp}(\boldsymbol{\lambda}),$$

and, for every $\boldsymbol{\lambda} \in M$ and every $\mathbf{s} \in \operatorname{supp}(\boldsymbol{\lambda})$, one has $\mathbf{s} > t_j$. Then define

$$\mathbf{M}_{k+1} := \mathcal{F}\big(\Delta(\mathbf{S}_{k+1})\big).$$

Hence $\mathbf{S}_{k+1}$ consists of finite timed consequence streams whose entries may be either elementary consequences or menu consequences from $\mathbf{M}_k$, subject to the requirement that every menu component is verified by the suffix that follows it. Finally, define the finite-depth envelopes

$$\mathbf{S}_\infty := \bigcup_{k \geq 0} \mathbf{S}_k, \qquad \mathbf{M}_\infty := \bigcup_{k \geq 0} \mathbf{M}_k.$$

Note that every element of $\mathbf{S}_\infty$ is a finite timed consequence stream, and every element of $\mathbf{M}_\infty$ is a finite menu of finitely supported lotteries. The union allows the finite nesting depth of menu consequences to be arbitrarily large across different finite objects.

Lemma 1. *The sequences $(\mathbf{S}_k)_{k \geq 0}$ and $(\mathbf{M}_k)_{k \geq 0}$ are increasing. Moreover,*

$$\mathbf{M}_\infty = \mathcal{F}\big(\Delta(\mathbf{S}_\infty)\big).$$

Finally, $\mathbf{S}_\infty$ is exactly the set of finite timed streams over $\mathbb{R} \times (Y \sqcup \mathbf{M}_\infty)$ satisfying the following self-verification condition: whenever a component $z_j = M \in \mathbf{M}_\infty$ occurs, one has $j < \tau$, there exists $\boldsymbol{\lambda} \in M$ such that $\mathbf{y}_{>j} \in \operatorname{supp}(\boldsymbol{\lambda})$, and, for every $\boldsymbol{\lambda} \in M$ and every $\mathbf{s} \in \operatorname{supp}(\boldsymbol{\lambda})$, one has $\mathbf{s} > t_j$.

Lemma 1 records three basic facts. First, monotonicity implies that allowing one more round of menu consequences never removes any streams or menus already available. Second, the limiting menu domain is closed under finite menus of finitely supported lotteries over streams in $\mathbf{S}_\infty$. Third, the limiting stream domain contains exactly the finite self-verifying streams: whenever a menu appears in a stream, the realised suffix must be one of the continuations encoded by that menu.

Proof. We first prove monotonicity. Since a primitive timed consequence chain contains no menu component, the self-verification requirement is vacuous on $\mathbf{S}_0$. Identifying Y with its image under the canonical injection into $Y \sqcup \mathbf{M}_0$, every timed consequence chain over $\mathbb{R} \times Y$ is therefore an element of $\mathbf{S}_1$. Hence $\mathbf{S}_0 \subseteq \mathbf{S}_1$. It follows that $\Delta(\mathbf{S}_0) \subseteq \Delta(\mathbf{S}_1)$, and therefore

$$\mathbf{M}_0 = \mathcal{F}(\Delta(\mathbf{S}_0)) \subseteq \mathcal{F}(\Delta(\mathbf{S}_1)) = \mathbf{M}_1.$$

Next, suppose $\mathbf{S}_k \subseteq \mathbf{S}_{k+1}$. Then $\Delta(\mathbf{S}_k) \subseteq \Delta(\mathbf{S}_{k+1})$, so

$$\mathbf{M}_k = \mathcal{F}(\Delta(\mathbf{S}_k)) \subseteq \mathcal{F}(\Delta(\mathbf{S}_{k+1})) = \mathbf{M}_{k+1}.$$

Take any $\mathbf{y} \in \mathbf{S}_{k+1}$. Since $\mathbf{M}_k \subseteq \mathbf{M}_{k+1}$, the stream $\mathbf{y}$ is a finite timed stream over $\mathbb{R} \times (Y \sqcup \mathbf{M}_{k+1})$. Moreover, every menu component of $\mathbf{y}$ already belongs to $\mathbf{M}_k$, and the witnesses that verify it as a member of $\mathbf{S}_{k+1}$ are unchanged when the ambient menu domain is enlarged to $\mathbf{M}_{k+1}$. Hence $\mathbf{y} \in \mathbf{S}_{k+2}$. Thus $\mathbf{S}_{k+1} \subseteq \mathbf{S}_{k+2}$. This proves monotonicity of both sequences.

We next prove the menu-domain identity. Monotonicity gives $\mathbf{M}_\infty \subseteq \mathcal{F}\big(\Delta(\mathbf{S}_\infty)\big)$, because for every k one has $\mathbf{M}_k = \mathcal{F}(\Delta(\mathbf{S}_k))$ and $\mathbf{S}_k \subseteq \mathbf{S}_\infty$. Conversely, let $F \in \mathcal{F}(\Delta(\mathbf{S}_\infty))$. Since F is finite and every lottery $\boldsymbol{\lambda} \in F$ has finite support, the set

$$A := \bigcup_{\boldsymbol{\lambda} \in F} \operatorname{supp}(\boldsymbol{\lambda})$$

is a finite subset of $\mathbf{S}_\infty$. For each $\mathbf{y} \in A$, choose $k(\mathbf{y})$ such that $\mathbf{y} \in \mathbf{S}_{k(\mathbf{y})}$, and set $K := \max_{\mathbf{y} \in A} k(\mathbf{y})$. By monotonicity, $A \subseteq \mathbf{S}_K$. Hence each $\boldsymbol{\lambda} \in F$ belongs to $\Delta(\mathbf{S}_K)$, and therefore

$$F \in \mathcal{F}\big(\Delta(\mathbf{S}_K)\big) = \mathbf{M}_K \subseteq \mathbf{M}_\infty.$$

This proves $\mathbf{M}_\infty = \mathcal{F}\big(\Delta(\mathbf{S}_\infty)\big)$.

It remains to identify $\mathbf{S}_\infty$. By construction, every $\mathbf{y} \in \mathbf{S}_\infty$ is a finite timed stream over $\mathbb{R} \times (Y \sqcup \mathbf{M}_\infty)$ satisfying the self-verification condition at every menu component. Conversely, suppose $\mathbf{y}$ is a finite timed stream over $\mathbb{R} \times (Y \sqcup \mathbf{M}_\infty)$ satisfying the self-verification condition. If $\mathbf{y}$ contains no menu component, then $\mathbf{y} \in \mathbf{S}_0 \subseteq \mathbf{S}_\infty$. Otherwise, since $\mathbf{y}$ is finite, it contains only finitely many menu components. For each menu component M appearing in $\mathbf{y}$, choose $k(M)$ such that $M \in \mathbf{M}_{k(M)}$, and let K be the maximum of these finitely many indices. By monotonicity, every menu component of $\mathbf{y}$ belongs to $\mathbf{M}_K$. The same witnesses that verify $\mathbf{y}$ relative to $\mathbf{M}_\infty$ verify it relative to $\mathbf{M}_K$. Hence $\mathbf{y} \in \mathbf{S}_{K+1} \subseteq \mathbf{S}_\infty$. □

Lemma 2. *Let $N \in \mathbf{M}_\infty$, and let $\mathbf{a} \in \mathbf{S}_\infty$. If*

$$\mathbf{a} \notin \bigcup_{\boldsymbol{\lambda} \in N} \operatorname{supp}(\boldsymbol{\lambda}),$$

then no stream in $\mathbf{S}_\infty$ has a component N whose suffix is exactly $\mathbf{a}$. In particular, suppressing dates, the two-stage stream $(N, \mathbf{a})$ does not belong to $\mathbf{S}_\infty$.

Proof. If such a stream belonged to $\mathbf{S}_\infty$, then the component N would have to satisfy the self-verification condition in Lemma 1. Hence the suffix following N, namely $\mathbf{a}$, would have to belong to $\operatorname{supp}(\boldsymbol{\lambda})$ for some $\boldsymbol{\lambda} \in N$. This contradicts the hypothesis. □

A.2 History-compatible tree domains

The enriched domains $(\mathbf{S}_\infty, \mathbf{M}_\infty)$ specify which complete consequence streams and menu consequences are admissible. We now impose the compatibility restrictions needed to evaluate a finite tree after a given history. Since $\mathbf{S}_\infty$ is the domain of complete self-verifying streams, a history need not itself belong to $\mathbf{S}_\infty$. A history may be a proper prefix of a complete self-verifying stream. Let H denote the set consisting of the empty history and all finite timed streams that are prefixes of some stream in $\mathbf{S}_\infty$.

For each $h \in H$, define

$$\bar{t}(h) := \max\Big\{t_j : (t_j, z_j) \text{ appears in } h\Big\},$$

with the convention $\bar{t}(\varnothing) = -\infty$. The number $\bar{t}(h)$ is the last date reached before the tree currently under consideration begins.

Given $h \in H$, a finite risky decision tree $T \in \mathcal{T}\big(\mathbb{R} \times (Y \sqcup \mathbf{M}_\infty)\big)$ is history-compatible after h if it satisfies the following two conditions. First, $t(n) > \bar{t}(h)$ for every timed consequence node $n \in N^y(T)$. Second, every finite stream obtained by concatenating h with the timed consequences encountered along an initial segment of a path in T belongs to H. For each history $h \in H$, define $\mathcal{T}_h$ as the class of all finite risky decision trees that are history-compatible after h. Thus $\mathcal{T}_h$ is the class of enriched decision trees that can be faced after history h without violating either chronology or the self-verification restrictions encoded in Lemma 1.

Given $T \in \mathcal{T}_h$ and $n \in N(T)$, let h_n denote the history obtained by concatenating h with all timed consequences encountered on the unique path from $n_0(T)$ to n, excluding any timed consequence at n itself. If no timed consequence node is encountered on this path, then $h_n = h$. By definition of $\mathcal{T}_h$, one has $h_n \in H$.

Lemma 3. *Let $h \in H$, $T \in \mathcal{T}_h$, and $n \in N(T)$. Then $T_{\geq n} \in \mathcal{T}_{h_n}$.*

The lemma says that history compatibility is preserved when one passes to a continuation subtree. Once T can be faced after h, the subtree $T_{\geq n}$ can be faced after the history accumulated by the time node n is reached.

Proof. Fix $h \in H$, $T \in \mathcal{T}_h$, and $n \in N(T)$. We first verify chronological compatibility. Take any timed consequence node $m \in N^y(T_{\geq n})$. We must show that $t(m) > \bar{t}(h_n)$. If $h_n = h$, then $\bar{t}(h_n) = \bar{t}(h)$. Since $T \in \mathcal{T}_h$, every timed consequence node of T, and therefore every timed consequence node of $T_{\geq n}$, has date strictly larger than $\bar{t}(h)$. Hence $t(m) > \bar{t}(h_n)$.

Suppose instead that $h_n \neq h$. Then h_n is obtained from h by appending the finite chain of timed consequences encountered on the unique path from $n_0(T)$ to n, excluding any timed consequence at n. Let r be the last timed consequence node on that path. Then $\bar{t}(h_n) = t(r)$. Since m belongs to the continuation subtree rooted at n, the unique path from $n_0(T)$ to m passes through n. Hence this path contains r before m, unless m is itself the first timed consequence node at or after n. In either case, the time-ordering of timed consequence chains in $\mathcal{T}\big(\mathbb{R} \times (Y \sqcup \mathbf{M}_\infty)\big)$ implies $t(m) > t(r) = \bar{t}(h_n)$. Therefore every timed consequence node of $T_{\geq n}$ occurs after $\bar{t}(h_n)$.

It remains to verify prefix compatibility. Take any finite stream obtained by concatenating h_n with the timed consequences encountered along an initial segment of a path in $T_{\geq n}$. By the definition of h_n, this stream is also obtained by concatenating h with the timed consequences encountered along an initial segment of the corresponding path in T. Since $T \in \mathcal{T}_h$, the resulting stream belongs to H. Hence $T_{\geq n}$ is prefix-compatible after h_n. We have shown both defining conditions of $\mathcal{T}_{h_n}$, so $T_{\geq n} \in \mathcal{T}_{h_n}$. $\square$

A.3 Feasible sets generated by a tree

We next define the feasible set generated by a finite enriched decision tree. For a tree T, let $\mathbf{P}(T)$ denote the set of path strategies in T. Given $\mathbf{p} \in \mathbf{P}(T)$, let $\boldsymbol{\eta}(T, \mathbf{p})$ denote the formal lottery over enriched consequence streams induced by following $\mathbf{p}$ through T. This lottery is constructed by the usual backward recursion (see H&TM). At a terminal timed consequence node it is the degenerate lottery on the one-element stream at that node. At an intermediate timed consequence node, the realised timed consequence is prepended to the lottery generated by the unique successor subtree. At a decision node, the recursion follows the successor selected by $\mathbf{p}$. At a chance node, it takes the mixture induced by the specified roulette lottery. The feasible set generated by T is

$$F(T) := \{\boldsymbol{\eta}(T, \mathbf{p}) : \mathbf{p} \in \mathbf{P}(T)\}.$$

For any node $n \in N(T)$, the feasible set $F(T_{\geq n})$ is defined in the same way for the continuation subtree rooted at n. This convention leaves the prior history outside $F(T)$. If T is evaluated after history h, and one wants the full realised chain including that history, one writes $h \circ \boldsymbol{\lambda}$ for $\boldsymbol{\lambda} \in F(T)$.

At this stage $F(T)$ is only a recursively defined feasible set. The next subsection imposes menu consistency. Proposition A.1 then shows that, for menu-consistent trees, the recursively generated feasible sets are finite menus of lotteries over the self-verifying stream domain $\mathbf{S}_\infty$.

A.4 Menu-consistent trees

The preceding subsection defined the feasible set generated by a finite enriched tree. We now impose the consistency requirement on trees with menu consequences. A timed consequence may belong to the menu domain $\mathbf{M}_\infty$, but in the intended interpretation a menu consequence records the feasible continuation opportunities generated by the subtree that follows it. We also impose the domain convention that menu consequences occur only at intermediate timed consequence nodes. Terminal timed consequence nodes carry primitive labels in $\mathbb{R} \times Y$.

Let $h \in H$, and let $T \in \mathcal{T}_h$. An intermediate timed consequence node $n \in N^i(T)$ is a menu-consequence node if its timed consequence is of the form $\boldsymbol{y}(n) = (t(n), M_n)$, where $M_n \in \mathbf{M}_\infty$. Since n is an intermediate node, it has a unique immediate successor, which we denote by $n^+ \in N^{+1}(n)$. We say that $T \in \mathcal{T}_h$ is *menu-consistent* just in case every terminal timed consequence node carries a primitive label, and, for every intermediate menu-consequence node $n \in N^i(T)$, one has $M_n = F(T_{\geq n^+})$. Define

$$\mathcal{T}_h^m := \Big\{ T \in \mathcal{T}_h : T \text{ is menu-consistent} \Big\}.$$

Thus $\mathcal{T}_h^m$ is the subclass of history-compatible trees in which every menu consequence is generated by the continuation subtree it is meant to describe. The equality $M_n = F(T_{\geq n^+})$ uses the recursively defined feasible set from Section A.3; requiring this equality is precisely the tree-level consistency restriction.

Lemma 4. *Let $h \in H$, $T \in \mathcal{T}_h^m$, and let $n \in N^i(T)$ be an intermediate menu-consequence node with unique successor n^+. If $\boldsymbol{y}(n) = (t(n), M_n)$, with $M_n \in \mathbf{M}_\infty$, then $M_n = F(T_{\geq n^+})$. In particular, in a menu-consistent tree, the menu consequence at n is uniquely determined by the continuation subtree that follows n.*

The lemma makes explicit the role of menu consistency. A menu-consequence node adds no branching, introduces no additional continuation opportunities, and removes none. It records the feasible set already generated by the successor subtree.

Proof. Because $T \in \mathcal{T}_h^m$, the tree is menu-consistent. By definition of menu consistency, every intermediate menu-consequence node $n \in N^i(T)$, with unique successor n^+ and timed consequence $\boldsymbol{y}(n) = (t(n), M_n)$, satisfies $M_n = F(T_{\geq n^+})$. This is the desired equality. Since $F(T_{\geq n^+})$ is determined by the successor subtree $T_{\geq n^+}$, the menu consequence at n is uniquely determined by that subtree. $\square$

Proposition A.1. *Let $h \in H$, $T \in \mathcal{T}_h^m$, and $n \in N(T)$. Then $F(T_{\geq n})$ is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$. Hence*

$$F(T_{\geq n}) \in \mathbf{M}_\infty.$$

The proposition connects the abstract construction of $(\mathbf{S}_\infty, \mathbf{M}_\infty)$ to menu-consistent finite decision trees. It shows that every feasible continuation set generated by such a tree is itself an admissible menu consequence.

Proof. Fix $h \in H$, $T \in \mathcal{T}_h^m$, and $n \in N(T)$. We prove the claim for every continuation subtree $T_{\geq n}$ by induction on the depth of $T_{\geq n}$. Since T is finite, this depth is finite.

If $n \in N^t(T)$ is a terminal timed consequence node, then, by menu consistency, the terminal label is primitive, say $\boldsymbol{y}(n) = (t(n), y(n))$ with $y(n) \in Y$. The only induced consequence stream is the one-element stream $(\boldsymbol{y}(n))$, which belongs to $\mathbf{S}_0 \subseteq \mathbf{S}_\infty$. Hence $F(T_{\geq n}) = \{\delta_{(\boldsymbol{y}(n))}\}$, which is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$.

Suppose next that $n \in N^i(T)$ is an intermediate timed consequence node with unique successor n^+. First consider the case in which $\boldsymbol{y}(n) = (t(n), y(n))$ with $y(n) \in Y$. By the induction hypothesis, $F(T_{\geq n^+})$ is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$. For each $\boldsymbol{\lambda} \in F(T_{\geq n^+})$, the lottery $\boldsymbol{y}(n) \circ \boldsymbol{\lambda}$ has finite support, and every stream in its support is obtained by prepending the primitive timed consequence $\boldsymbol{y}(n)$ to a stream in $\mathbf{S}_\infty$. This introduces no new menu component, so all existing menu components remain self-verifying. Moreover, by Lemma 3, the successor subtree $T_{\geq n^+}$ is chronologically compatible after the history that includes $\boldsymbol{y}(n)$. Hence every continuation stream generated by $T_{\geq n^+}$ has all its dates strictly greater than $t(n)$. Therefore the resulting streams belong to $\mathbf{S}_\infty$, and

$$F(T_{\geq n}) = \{\boldsymbol{y}(n) \circ \boldsymbol{\lambda} : \boldsymbol{\lambda} \in F(T_{\geq n^+})\}$$

is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$.

Now consider the case in which n is an intermediate menu-consequence node, so that $\boldsymbol{y}(n) = (t(n), M_n)$ with $M_n \in \mathbf{M}_\infty$. By Lemma 4,

$$M_n = F(T_{\geq n^+}).$$

By the induction hypothesis, $F(T_{\geq n^+})$ is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$. Take any $\boldsymbol{\lambda} \in F(T_{\geq n^+}) = M_n$, and any $\mathbf{s} \in \operatorname{supp}(\boldsymbol{\lambda})$. The stream obtained by prepending

$(t(n), M_n)$ to $\mathbf{s}$ verifies the new menu component, because its suffix after M_n is exactly $\mathbf{s}$, and $\mathbf{s} \in \operatorname{supp}(\boldsymbol{\lambda})$ for some $\boldsymbol{\lambda} \in M_n$. Existing menu components inside $\mathbf{s}$ remain self-verifying. Moreover, by Lemma 3, the successor subtree $T_{\geq n^+}$ is chronologically compatible after the history that includes $(t(n), M_n)$. Hence every stream in the support of every lottery in $F(T_{\geq n^+}) = M_n$ has all its dates strictly greater than $t(n)$. The date requirement in Lemma 1 is therefore satisfied. Thus $\boldsymbol{y}(n) \circ \boldsymbol{\lambda} \in \Delta(\mathbf{S}_\infty)$ for every $\boldsymbol{\lambda} \in F(T_{\geq n^+})$, and so

$$F(T_{\geq n}) = \{\boldsymbol{y}(n) \circ \boldsymbol{\lambda} : \boldsymbol{\lambda} \in F(T_{\geq n^+})\}$$

is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$.

Suppose that $n \in N^d(T)$ is a decision node. For each $m \in N^{+1}(n)$, the continuation subtree $T_{\geq m}$ has smaller depth, so the induction hypothesis implies that $F(T_{\geq m})$ is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$. Since $N^{+1}(n)$ is finite and non-empty, the union

$$F(T_{\geq n}) = \bigcup_{m \in N^{+1}(n)} F(T_{\geq m})$$

is again a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$.

Finally, suppose that $n \in N^c(T)$ is a chance node with roulette lottery $\pi(\cdot \mid n)$ on $N^{+1}(n)$. For each $m \in N^{+1}(n)$, the induction hypothesis implies that $F(T_{\geq m})$ is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$. Hence

$$F(T_{\geq n}) = \left\{ \sum_{m \in N^{+1}(n)} \pi(m \mid n) \boldsymbol{\lambda}_m : \boldsymbol{\lambda}_m \in F(T_{\geq m}) \text{ for every } m \in N^{+1}(n) \right\}.$$

Because $N^{+1}(n)$ and each $F(T_{\geq m})$ are finite and non-empty, this set is finite and non-empty. Each element is a finite convex combination of finitely supported lotteries over $\mathbf{S}_\infty$, and is therefore itself an element of $\Delta(\mathbf{S}_\infty)$. Thus $F(T_{\geq n})$ is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$.

The induction proves that $F(T_{\geq n}) \in \mathcal{F}(\Delta(\mathbf{S}_\infty))$. By Lemma 1, $\mathcal{F}(\Delta(\mathbf{S}_\infty)) = \mathbf{M}_\infty$. Hence $F(T_{\geq n}) \in \mathbf{M}_\infty$. □

A.5 The history-indexed domain of menu-consequence trees

We can now assemble the domain used in the paper. The unrestricted enriched tree domain is $\mathcal{T}\big(\mathbb{R} \times (Y \sqcup \mathbf{M}_\infty)\big)$, the syntactic class of finite risky decision trees with decision nodes, chance nodes, and timed consequence nodes whose timed consequences lie in $\mathbb{R} \times (Y \sqcup \mathbf{M}_\infty)$. Thus each timed consequence is either an elementary timed consequence (t, y) with $y \in Y$, or a timed menu consequence (t, M), with $M \in \mathbf{M}_\infty$. As in the baseline timed-consequence model, each intermediate timed consequence node has a unique immediate successor.

The unrestricted class is larger than the domain relevant for behaviour. A tree must be history-compatible with the history after which it is faced, and its menu consequences must be menu-consistent. In particular, within the behaviourally relevant domain, terminal timed consequence nodes carry primitive labels, while menu consequences occur only at intermediate timed consequence nodes and record the feasible set generated by the successor subtree. We therefore define the history-indexed domain of menu-consequence trees by

$$\mathcal{D}^m := \Big\{(h, T) : h \in H \text{ and } T \in \mathcal{T}_h^m\Big\}.$$

Equivalently, $\mathcal{D}^m$ is the class of all history–tree pairs for which the tree can be faced after the history and every intermediate menu consequence records the feasible set generated by the continuation subtree that follows it.

Proposition A.2. *Let $(h, T) \in \mathcal{D}^m$. Then, for every node $n \in N(T)$, one has $T_{\geq n} \in \mathcal{T}_{h_n}^m$. Moreover $F(T_{\geq n}) \in \mathbf{M}_\infty$, and for every path strategy $\mathbf{p} \in \mathbf{P}(T_{\geq n})$, one has $\boldsymbol{\eta}(T_{\geq n}, \mathbf{p}) \in \Delta(\mathbf{S}_\infty)$. In particular, $(h_n, T_{\geq n}) \in \mathcal{D}^m$.*

This proposition records the closure property needed for dynamic choice. Starting from an admissible pair (h, T), every continuation subtree is again an admissible tree after the history accumulated at its root. Feasible continuation sets remain admissible menu consequences, and every induced consequence lottery is a lottery over self-verifying streams in $\mathbf{S}_\infty$.

Proof. Fix $(h, T) \in \mathcal{D}^m$ and $n \in N(T)$. Since $(h, T) \in \mathcal{D}^m$, we have $T \in \mathcal{T}_h^m$, and hence $T \in \mathcal{T}_h$. By Lemma 3, $T_{\geq n} \in \mathcal{T}_{h_n}$.

It remains to show that $T_{\geq n}$ is menu-consistent. First, let $m \in N^t(T_{\geq n})$ be a terminal timed consequence node of $T_{\geq n}$. Since $T_{\geq n}$ is a continuation subtree of T, the node m is also a terminal timed consequence node of T. Because $T \in \mathcal{T}_h^m$, terminal timed consequence nodes in T carry primitive labels. Hence the same is true of terminal timed consequence nodes in $T_{\geq n}$.

Second, let $m \in N^i(T_{\geq n})$ be an intermediate menu-consequence node of $T_{\geq n}$, and let m^+ be its unique successor. Since $T_{\geq n}$ is a continuation subtree of T, the node m is also an intermediate menu-consequence node of T, with the same unique successor m^+. Because $T \in \mathcal{T}_h^m$, menu consistency of T implies that, if $\boldsymbol{y}(m) = (t(m), M_m)$, then $M_m = F(T_{\geq m^+})$. The continuation subtree of $T_{\geq n}$ rooted at m^+ is exactly $T_{\geq m^+}$. Therefore the same equality gives the menu-consistency condition for $T_{\geq n}$. Hence $T_{\geq n} \in \mathcal{T}_{h_n}^m$.

Since $T_{\geq n} \in \mathcal{T}_{h_n}^m$, Proposition A.1 implies that $F(T_{\geq n})$ is a finite non-empty subset of $\Delta(\mathbf{S}_\infty)$, and hence $F(T_{\geq n}) \in \mathbf{M}_\infty$. Finally, for any $\mathbf{p} \in \mathbf{P}(T_{\geq n})$, the induced lottery

$\boldsymbol{\eta}(T_{\geq n}, \mathbf{p})$ belongs to $F(T_{\geq n})$ by definition of the feasible set. Since $F(T_{\geq n}) \subseteq \Delta(\mathbf{S}_\infty)$, it follows that $\boldsymbol{\eta}(T_{\geq n}, \mathbf{p}) \in \Delta(\mathbf{S}_\infty)$. Thus $(h_n, T_{\geq n}) \in \mathcal{D}^m$, as claimed. □

The construction therefore yields the domain used in the main text. The stream domain $\mathbf{S}_\infty$ contains exactly the finite self-verifying streams, and $\mathbf{M}_\infty = \mathcal{F}(\Delta(\mathbf{S}_\infty))$ is the associated domain of finite menus of lotteries over such streams. Histories are admissible prefixes of elements of $\mathbf{S}_\infty$, while $\mathcal{T}_h$ is the class of trees that can be faced after history h without violating chronology or prefix admissibility. The subclass $\mathcal{T}_h^m$ imposes the additional requirement that every menu consequence in the tree truthfully records the feasible set generated by its successor subtree.

Propositions A.1 and A.2 are the two closure results needed later. The first shows that every menu-consistent tree generates feasible sets in $\mathbf{M}_\infty$, equivalently lotteries over self-verifying streams in $\mathbf{S}_\infty$. The second shows that this property is stable under passage to continuation subtrees: if $(h, T) \in \mathcal{D}^m$, then every continuation pair $(h_n, T_{\geq n})$ also belongs to $\mathcal{D}^m$. Thus menu-consistent dynamic trees generate exactly the kind of stream lotteries on which the expected-utility representation in the main text is defined.

Appendix B Proof of Proposition 1

We first record the normal-form comparison step used in the proof. The comparison trees used here are normal-form trees over already constructed streams in $\mathbf{S}_\infty$. They are not dynamic menu-consequence trees in $\mathcal{T}_h^m$, and they do not re-expand the menu components of a stream into a new dynamic tree.

Lemma 5. *For every $\boldsymbol{\lambda}, \boldsymbol{\mu} \in \Delta(\mathbf{S}_\infty)$, there exists a normal-form comparison tree over stream consequences whose feasible set is $\{\boldsymbol{\lambda}, \boldsymbol{\mu}\}$.*

Proof. Construct a tree whose initial node is a decision node with two immediate successors. The first branch leads to a chance node resolving the finitely supported lottery $\boldsymbol{\lambda}$ over terminal stream consequences in $\mathbf{S}_\infty$. The second branch leads to a chance node resolving the finitely supported lottery $\boldsymbol{\mu}$ over terminal stream consequences in $\mathbf{S}_\infty$. If either lottery is degenerate, the corresponding chance node may be replaced by a terminal stream-consequence node. The feasible set of the resulting normal-form tree is exactly $\{\boldsymbol{\lambda}, \boldsymbol{\mu}\}$. Its terminal objects are elements of $\mathbf{S}_\infty$. The tree contains no internal menu-consequence node, so menu consistency is vacuous. □

Fix $h \in H$, and take any $\boldsymbol{\lambda}, \boldsymbol{\mu} \in \Delta(\mathbf{S}_\infty)$. By Lemma 5, there exists a normal-form comparison tree $T^{\boldsymbol{\lambda},\boldsymbol{\mu}}$ over stream consequences such that $F(T^{\boldsymbol{\lambda},\boldsymbol{\mu}}) = \{\boldsymbol{\lambda}, \boldsymbol{\mu}\}$. Since $\succsim_h$ is prerational, there exists a consequentialist normal-form invariant behaviour rule

β such that the induced base preference $\succsim_h^\beta$ coincides with $\succsim_h$. By consequentialist normal-form invariance, the behaviour rule induces a history-indexed choice function $\mathcal{C}_h^\beta$ on finite feasible sets of lotteries over $\mathbf{S}_\infty$. In particular,

$$\Phi_h^\beta(T^{\boldsymbol{\lambda},\boldsymbol{\mu}}) = \mathcal{C}_h^\beta(\{\boldsymbol{\lambda},\boldsymbol{\mu}\}).$$

The behaviour rule is non-empty-valued, so $\mathcal{C}_h^\beta(\{\boldsymbol{\lambda},\boldsymbol{\mu}\})$ is a non-empty subset of $\{\boldsymbol{\lambda},\boldsymbol{\mu}\}$. Hence either $\boldsymbol{\lambda} \in \mathcal{C}_h^\beta(\{\boldsymbol{\lambda},\boldsymbol{\mu}\})$ or $\boldsymbol{\mu} \in \mathcal{C}_h^\beta(\{\boldsymbol{\lambda},\boldsymbol{\mu}\})$. By definition of the induced base preference, either $\boldsymbol{\lambda} \succsim_h^\beta \boldsymbol{\mu}$ or $\boldsymbol{\mu} \succsim_h^\beta \boldsymbol{\lambda}$. Since $\succsim_h^\beta = \succsim_h$, either $\boldsymbol{\lambda} \succsim_h \boldsymbol{\mu}$ or $\boldsymbol{\mu} \succsim_h \boldsymbol{\lambda}$. This proves that $\succsim_h$ is complete. $\square$

The proof uses normal-form comparison over streams. It should therefore be read together with the history-indexed domain constructed in Appendix A. If a menu consequence $F = \{a, b, c\}$ has already occurred, then F belongs to the history at which subsequent behaviour is evaluated. Thus, after history $h = (F)$, a continuation tree may offer a choice between b and c, generating the full streams $h \circ b = (F, b)$ and $h \circ c = (F, c)$. This recovers the comparison between the two unchosen alternatives without holding fixed a false future menu. The menu consequence is fixed because it is remembered as part of the past.

A different issue arises when the streams to be compared were generated under different menus. For example, let $M = \{a, b\}$ and $N = \{c, d\}$. The streams (M, a) and (N, c) are self-verifying, but a dynamic tree that first records M and then offers only a, or first records N and then offers only c, would not be menu-consistent. The binary comparison in the proof above is not such a dynamic tree. It is a normal-form comparison tree whose terminal consequences are the already constructed streams (M, a) and (N, c). Hence no new menu label is introduced, and no menu label is falsified. Menu consistency constrains dynamic menu-consequence trees; it does not require normal-form comparison trees over already constructed streams to re-expand those streams into new dynamic menu trees.

References

Adler, M.A. (2025) “Prioritarianism as a Theory of Value” in *Stanford Encyclopedia of Philosophy.*

Alkire, Sabina (2008) “Concepts and measures of agency” in *Arguments for a Better World: Essays in Honor of Amartya Sen, Part IV: Identity, Collective Action and Public Economics* Oxford University Press, ch. 24.

Anscombe, Frank, and Aumann, Robert J. (1963) “A Definition of Subjective Probability”. *Annals of Mathematical Statistics* 34 (1): 199–205.

Arrow, Kenneth J., and Anthony C. Fisher (1974) “Environmental Preservation, Uncertainty, and Irreversibility” *Quarterly Journal of Economics* 88(2): 312–319.

Atkinson, A. B. (1970) “On the measurement of inequality” *Journal of Economic Theory* 2 (3): 244–263.

Barberà, Salvador, Walter Bossert, and Prasanta K. Pattanaik (2004) “Ranking Sets of Objects” in *Handbook of Utility Theory, Vol. 2: Extensions* ch. 17, pp. 893–977.

Bergson, Abram (1938) “A Reformulation of Certain Aspects of Welfare Economics” *Quarterly Journal of Economics* 52 (2): 310–334.

Brown, Campbell (2011) “Consequentialize This” *Ethics* 121 (4): 749–771.

Chernoff, H (1954) Rational selection of decision functions *Econometrica* 22:422-443

Dekel, E., Lipman, B. L., and Rustichini, A. (2001). Representing preferences with a unique subjective state space. *Econometrica*, 69 (4): 891–934.

Dekel, E., Lipman, B. L., Rustichini, A., and Sarver, T. (2007). Representing preferences with a unique subjective state space: A corrigendum *Econometrica*, 75 (2): 591–600.

Diamond, Peter A. (1967) “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment” *Journal of Political Economy* 75 (5): 765–766.

Dillenberger, D. and Sadowski, P. (2012). Ashamed to be selfish. Theoretical Economics, 7 (1): 99–124.

Drèze, Jacques H. (ed.) (1974) *Allocation Under Uncertainty: Equilibrium and Optimality* (London, Macmillan).

Drèze, Jacques H. (1987) *Essays on economic decisions under uncertainty* (Cambridge, Cambridge University Press)

Drèze, Jean and Sen, Amartya K. (2002). *India, development and participation, 2nd ed.* New Delhi, New York: Oxford University Press.

Ellsberg, Daniel (1961) "Risk, Ambiguity, and the Savage Axioms" *Quarterly Journal of Economics* 75 (4): 643–669.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57 (4): 937–969.

Gaertner, Wulf, Prasanta Pattanaik, and Kotaro Suzumura (1992) "Individual Rights Revisited" *Economica* 59: 161–77.

Gärdenfors, P., and Sahlin, N.E. (1982) "Unreliable probabilities, risk taking, and decision making". *Synthese* 53: 361–386.

Gibbard, Allan (1974) "A Pareto-consistent libertarian claim". *Journal of Economic Theory* 7: 388–410.

Gul, F. and Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, 69 (6): 1403–1435.

Hammond, Peter J. (1975) "A Note on Extreme Inequality Aversion" *Journal of Economic Theory* 11: 465–467.

Hammond, P. J. (1976). Changing Tastes and Coherent Dynamic Choice. *Review of Economic Studies* 43 (1): 159–173.

Hammond, P. J. (1988). "Consequentialist foundations for expected utility" *Theory and decision*, 25: 25–78.

Hammond, P.J. (1995) "Social Choice of Individual and Group Rights" in W.A. Barnett, H. Moulin, M. Salles, and N. Schofield (eds.) *Social Choice, Welfare, and Ethics* (CUP), ch. 3, pp. 55–77.

Hammond, P.J. (1997) "Game Forms versus Social Choice Rules as Models of Rights" in K.J. Arrow, A.K. Sen, and K. Suzumura (eds.) *Social Choice Re-examined, Vol. II* (Macmillan) ch. 11, pp. 82–95.

Hammond, P. J. (1998). Objective expected utility: A consequentialist perspective. *Handbook of utility theory, volume 1* (Springer) 143–211.

Hammond, P. J. (2022). Prerationality as avoiding predictably regrettable consequences. *Revue économique*, 73 (6): 943–976.

Hammond, P.J. (2026) "Bounded Rationality with Subjective Evaluations in Enlivened but Truncated Decision Trees" *Theory and Decision* https://doi.org/10.1007/s11238-025-10119-y.

Hammond, P. J. and Troccoli-Moretti, A. (2026). Prerationality in finite decision trees with non-terminal consequences Working Paper.

Harsanyi, John C. (1953) "Cardinal Utility in Welfare Economics and in the Theory of Risk-taking" *Journal of Political Economy* 61: 434–435.

Harsanyi, John C. (1955) "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility" *Journal of Political Economy* 63: 309–321.

Hart, A. G. (1940). *Anticipations, uncertainty, and dynamic planning.* University of Chicago Press, Chicago.

Huber, J., Payne, J. W., and Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of consumer research*, 9 (1): 90–98.

Keynes, J. Maynard (1930) *Treatise on Money* (London: Macmillan).

Keynes, J. Maynard (1936) *General Theory of Employment, Interest, and Money* (London: Macmillan).

Koopmans, T. C. (1950). Utility analysis of decisions affecting future well-being. (abstract). *Econometrica*, 18 (2): 174–175.

Koopmans, T. C. (1964). On the flexibility of future preferences. In *Human Judgments and Optimality*, ed. by M. W. Shelly and G. L. Bryan (New York: John Wiley).

Kopylov, I. and Noor, J. (2018). Commitments and weak resolve. *Economic Theory*, 66 (1): 1–19.

Kreps, D. M. (1979). A representation theorem for "preference for flexibility". *Econometrica*, pages 565–577.

Kreps, D. M. and Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46 (1): 185–200.

Lipman, B. L. and Pesendorfer, W. (2013). Temptation. In *Advances in economics and econometrics: Tenth World Congress, volume 1*, (Cambridge: Cambridge University Press) pp. 243–288.

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27 (4): 1622–1668.

Moulin, Hervé, and Peleg, Bezalel (1982) "Cores of effectivity functions and implementation theory" *Journal of Mathematical Economics* 10 (1): 115–145.

Noor, J. and Ren, L. (2023). Temptation and guilt. *Games and Economic Behavior*, 140:272–295.

Nussbaum, Martha (2000) *Women and Human Development: The Capabilities Approach* (Cambridge: Cambridge University Press).

Nussbaum, Martha (2011) *Creating Capabilities: The Human Development Approach* (Cambridge, MA: Harvard University Press).

Portmore, Douglas W. (2007) "Consequentializing Moral Theories" *Pacific Philosophical Quarterly* 88(1): 39–73.

Portmore, Douglas W. (2022) "Consequentializing" *Stanford Encyclopedia of Philosophy.*

Portmore, Douglas W. (2023) "Consequentializing agent-centered restrictions: A Kantsequentialist approach" *Analytic Philosophy* 64: 443–467.

Rabinowicz, Wlodek (2002) "Prioritarianism for Prospects" *Utilitas* 14(1): 2–21.

Rawls, John (1971) *A Theory of Justice* (Cambridge, MA: Harvard University Press).

Robeyns, Ingrid and Morten Fibieger Byskov (2020) "The Capability Approach" *Stanford Encyclopedia of Philosophy.*

Rowley, C. K. and Peacock, A. T. (1975) *Welfare Economics: A Liberal Restatement.*

Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica*, 76 (2): 263–305.

Savage, L.J. (1954) *The Foundations of Statistics* (New York: John Wiley).

Sen, A.K., (1970a) *Collective Choice and Social Welfare* (San Francisco: Holden-Day)

Sen, Amartya K. (1970b) "The Impossibility of a Paretian Liberal" *Journal of Political Economy* 78(1): 152–157.

Sen, A.K., 1985a. "Well-being, agency and freedom: The Dewey lectures 1984." *Journal of Philosophy* 82(4), p.169-221.

Sen, A.K., 1985b. Commodities and capabilities. (Amsterdam; New York: North-Holland; Elsevier Science).

Sen, A.K., 1993. Capability and Well-being. In Nussbaum, M. & Sen, A. (eds.) *The Quality of Life* (Oxford: Clarendon Press. p. 30-53.

Sen, A. (1997). Maximization and the act of choice. *Econometrica*, pages 745–779.

Sen, A.K., 1999. *Development as Freedom. 1st ed.* New York: Knopf Press.

Sen, A.K., 2000. Consequential Evaluation and Practical Reason. *Journal of Philosophy*, 97(9): 477–502

Sen, A.K., 2002. *Rationality and freedom.* Cambridge, Mass: Belknap Press.

Sugden, Robert (1986, revised 2005) *The Economics of Rights, Co-operation and Welfare* (Blackwell; then Palgrave Macmillan).

Sugden, Robert (2004) "The Opportunity Criterion: Consumer Sovereignty without the Assumption of Coherent Preferences" *American Economic Review* 94 (4): 1014–1033.

Sugden, Robert (2018) *The Community of Advantage: A Behavioural Economist's Defence of the Market* (Oxford University Press).

Temkin, Larry S. (2003) "Equality, priority or what?" *Economics and Philosophy* 19 (1): 61–87

Troccoli Moretti, A. (2024). Disappointment, risk aversion and dynamic depletion of self-control. Working Paper.

Vickrey, William S. (1945) "Measuring Marginal Utility by Reactions to Risk" *Econometrica* 13 (4): 319–333.

Welzel, Christian, and Ronald Inglehart (2010) "Agency, Values, and Well-Being: A Human Development Model" *Social Indicators Research* 97: 43–63