# Discounting Desirable Gambles

**Gregory Wheeler** G.WHEELER@FS.DE
*Frankfurt School of Finance & Management, Germany*

## Abstract

The desirable gambles framework offers the most comprehensive foundations for the theory of lower previsions, which in turn affords the most general account of imprecise probabilities. Nevertheless, for all its generality, the theory of lower previsions rests on the notion of linear utility. This commitment to linearity is clearest in the coherence axioms for sets of desirable gambles. This paper considers two routes to relaxing this commitment. The first preserves the additive structure of the desirable gambles framework and the machinery for coherent inference but detaches the interpretation of desirability from the multiplicative scale invariance axiom. The second strays from the additive combination axiom to accommodate repeated gambles that return rewards by a non-stationary processes that is not necessarily additive. Unlike the first approach, which is a conservative amendment to the desirable gambles framework, the second is a radical departure. Yet, common to both is a method for describing rewards called *discounted utility*.

**Keywords:** desirability, ergodicity, non-linear utility, lower previsions, imprecise probability

## 1. Lower Previsions: A Potted History

*In our opinion they are not fit for characterizing a new, weaker kind of coherent behavior.*

– De Finetti and Savage, 1962

A brief overview of the theory of lower previsions will help motivate my proposals for discounting desirable gambles and make the paper self-contained. But a short history is worth recounting on its own, least one mistake de Finetti and Savage's response to [27] as the last word on lower previsions rather than among the first.

Cedric Smith, with his proposal for inference with lower and upper "pignic odds" [27], showed that every lower prevision may be understood as the lower envelope of some set of linear previsions, an idea that Walley later generalized to characterize coherent lower previsions for bounded gambles. A *bounded gamble*, to fix notation, is a bounded real map $f$ from a set of possibilities $\Omega$, $f : \Omega \to \mathbb{R}$, interpreted as a gain or loss associated with each state $\omega \in \Omega$. When the possibility space $\Omega$ is fixed or clear from context, I simply write $f$ instead of '$f(\omega)$, for all $\omega \in \Omega$', and write $\mathbb{G}$ instead of '$\mathbb{G}(\Omega)$' to refer to the set of all gambles on $\Omega$. Binary arithmetic operations are point-wise operations over states. So, for example, $f - g$ is short for '$f(\omega) - g(\omega)$, for all $\omega \in \Omega$'. Gambles are random variables, so the Smith-Walley approach to defining lower previsions on gambles may be viewed (in hindsight) as an extension of de Finetti's theory of linear previsions [8].

Returning to Walley's characterization, a lower prevision $\underline{\mathrm{P}}$ defined on a class $\mathbb{C} \subseteq \mathbb{G}$ of gambles *avoids sure loss* if and only if there is a linear prevision P such that $\mathrm{P}(f) \geq \underline{\mathrm{P}}(f)$, for all gambles in $\mathbb{C}$; and $\underline{\mathrm{P}}$ is a *coherent lower prevision* if and only if there is a set of linear previsions $\mathbb{P}$ such that $\underline{\mathrm{P}}$ is the lower envelope of $\mathbb{P}$, that is, such that $\underline{\mathrm{P}}(f) = \inf\{\mathrm{P}(f) : \mathrm{P} \in \mathbb{P}\}$, for all $f$ in $\mathbb{C}$ [31, §3.3]. When the range of gambles $f$ in a class is restricted to $\{0,1\}$, then each $f$ is a characteristic function for a subset of $\Omega$, denoting an event, and $\underline{\mathrm{P}}(f)$ expresses a lower probability.

Walley's results connect together two models of imprecise probabilities: one that defines the functional $\underline{\mathrm{P}}$ as simply the lower expectation values of a given set of closed convex probability mass functions $\mathbb{P}$, called a *credal set*, and another which starts with coherence axioms for a functional $\underline{\mathrm{P}}$ that entail there is *some* credal set whose lower expectation values for each bounded gamble $f$ witness $\underline{\mathrm{P}}(f)$. Since de Finetti's theory of linear previsions falls out as a special case, Walley's unified theory is better understood as a more general characterization of coherent behavior rather than an entirely new and weaker one.

Sure loss avoidance and coherence extend to conditional lower previsions too [31], but with a catch, as inference in the credal set framework is forbiddingly complicated [5, 19, 28]. An alternative model for coherent lower previsions, constructed in terms of sets of *desirable gambles* [38], encompasses an even broader range of models for coherent behavior, one that strictly includes the credal set framework [32, 35] while greatly simplifying inference [29]. So, whereas the original credal set approach is restricted to bounded gambles, ties coherence criteria to sets of probabilities, but is impractical for predictive inference, the desirable gambles approach puts coherence and inference first, extends easily to include unbounded gambles, but demotes sets of probabilities to a secondary role, to be derived when possible or disposed of as necessary [36].

Nevertheless, for all its generality, the theory of lower previsions appears to rest on the concept of linear utility, a point made explicit by the coherence axioms for desirable gambles (Section 2). Apart from mathematical convenience, which should not be mistaken for a normative principle, the case for linear utility is thin. The principal arguments in favor of linear utility either concede that the marginal value of wealth decreases, by stipulating that linear utility is only warranted when small (but not too small!) stakes are at risk [8], or offer no argument but instead simply restate the linearity condition by positing an exogenous probability on which a "currency" of hypothetical lotteries is constructed [31, §2.2]. Alternatively, some retain the underlying mechanics of de Finetti's coherence criteria but abandon the *normative* commitment to linear utility. The fundamental theorem of "no arbitrage" asset pricing takes this tack, which is simply de Finetti's fundamental theorem of prevision but with discounting [3].

A weak case for linear utility and good reasons to favor non-linear utility (Section 3) raise the question of whether sets of desirable gambles can accommodate discounting. This paper introduces two approaches to discounting gambles, and from those two proposals follow two types of answers. The first (Section 4) prioritizes preserving the underlying coherence machinery of the desirable gambles framework and therefore may be viewed as an analog to the adoption of de Finetti's ideas by mathematical finance in the 1970s. The key idea is to preserve the additive structure of desirable gambles but separate the mathematical role non-negative multiplicative scaling plays from the behavioral and normative properties of desirability. For this new task, an adjustable utility scale is introduced to discount gambles. The second approach (Section 6) uses the same discounted utility but drops the additive structure of gambles, which allows for the introduction of non-additive dynamics in sequential gambles. This approach is a departure from the sets of desirable gambles framework, however. For although the additive dynamics assumed by the sets of desirable gambles approach appear as a special case, a type of lower expectation appears in this setting involving uncertainty over the type of dynamics rather than from uncertainty over state-dependent outcomes. However, the same adjustable discounted utility is used in both accounts, which includes linear utility as a special case.

The notion that desire satisfaction is concave rather than linear is as old as the concept of utility itself [1]. The aim of this paper instead is to introduce discounted utility in a form that is compatible with the theory of lower previsions, thus extending its reach while preserving its practical advantages, but to also show how the notion of discounted utility can be used to investigate imprecision that can arise from uncertain dynamics.

## 2. Desirable Gambles and Linear Utility

I mentioned in the introduction that the credal set approach formulates coherence axioms for a lower prevision functional $\underline{P}$ on a domain of bounded gambles. The desirable gambles framework, by contrast, consists of axioms for constructing a coherent set $\mathbb{D} \subseteq \mathbb{G}$ of bounded gambles directly from an initial, possibly empty, collection of gambles. For instance, a decision maker might be disposed to accept a finite collection of gambles and the axioms then may be used to determine whether that disposition is coherent, and therefore whether committing to those gambles would avoid sure loss. Following [29], there are four core coherence axioms for desirable gambles. For bounded gambles $f, g$ and positive real number $\lambda \geq 0$, a set of bounded gambles $\mathbb{D}$ is coherent when the following four conditions are satisfied:

A1. If $f < 0$, then $f \notin \mathbb{D}$ **(Avoid partial loss)**

A2. If $f \geq 0$, then $f \in \mathbb{D}$ **(Accept partial gain)**

A3. If $f \in \mathbb{D}$, then $\lambda f \in \mathbb{D}$ **(Positive Scale Invariance)**

A4. If $f \in \mathbb{D}$ and $g \in \mathbb{D}$, then $f + g \in \mathbb{D}$ **(Combination)**

Axioms A1 and A2 are rationality conditions: positive payments are desirable (A2) while negative payments are not (A1). Axiom A3 says that the desirability of a gamble is unchanged by the introduction of a positive scale and axiom A4 says that desirability is additive.[1]

Although rationality arguments are given for A3 and A4 [31], the axioms are principally closure operations. Axiom A3 ensures that a set $\mathbb{D}$ is a *cone*, since for every $f$ in $\mathbb{D}$ and $\lambda \geq 0$, $\lambda f \in \mathbb{D}$; and $\mathbb{D}$ is a *convex cone* if it is closed under both non-negative multiplication (A3) and addition (A4); that is, if, for any $f_1$ and $f_2$ in $\mathbb{D}$ and $\lambda_1, \lambda_2 \geq 0$, then $\lambda_1 f_1 + \lambda_2 f_2 \in \mathbb{D}$. A point of the form $\lambda_1 f_1 + \cdots + \lambda_n f_n$ with $\lambda_1, \ldots, \lambda_n \geq 0$ is a *conic combination* of $f_1, \ldots, f_n$, and a convex cone of $\mathbb{D}$ contains all conic combinations of its elements. The *conic hull* of $\mathbb{D}$ is the smallest convex cone,

$$\mathsf{cone}(\mathbb{D}) := \left\{ \sum_{i=1}^{n} \lambda_i f_i : f \in \mathbb{D}, \lambda_i \geq 0 \right\} \tag{1}$$

which contains $\mathbb{D}$. Thus, a coherent set of desirable gambles over a space of $|\Omega| = m$ possibilities is the convex hull in $\mathbb{R}^m$ that includes the non-negative orthant in $\mathbb{R}^m$, by A2, and excludes the negative orthant, by A1.[2]

Axioms A3 and A4 together encode the assumption that the utility of coherent gambles is measured on a linear scale. The reward $x \in \mathbb{R}$ of a gamble $f$ depends on which state $\omega$ obtains, so a gamble $f$ over $|\Omega| = m$ possible states describes a vector $\mathbf{x}_f = (x_1, \ldots, x_m)$ of state-contingent outcomes in $\mathbb{R}^m$. A *utility* is a real-valued function on $\mathbb{D}$ that

1. One finds the terms 'acceptable' and 'desirable' used interchangeably by some, but not all, authors. The main issue criss-crossing the literature is whether the status-quo 0-gamble is included in the set of coherent gambles. The approach I favor is to let the axioms settle such questions, and use the terms 'acceptable' and 'desirable' very loosely when speaking, as I am here, of the general theory.
2. When gambles are unbounded, coherence is defined with respect to a particular cone [6] [29, §13.2].

preserves the ordering of $\mathbf{x}_f = (x_1, \ldots, x_m)$ up to a constant $c$. From Equation 1, the inner product expression of expected utility is recovered immediately by setting all weights $w_i$ to one, with an $m$-vector $\mathbf{1}$ of unit weights, and letting the constant $c$ be zero:

$$u_{lin}(\mathbf{x}) = \left(\sum_{i=1}^{m} w_i x_i\right) + c \tag{2}$$
$$= \mathbf{1x} + 0$$
$$= \mathbf{x}$$

With these preliminaries in place, I turn next to introduce a discounted utility function that has one adjustable parameter whose range includes linear utility, as a corner case, but is otherwise a concave function of a reward $x$.

## 3. Discounted Utility

The original notion of utility, introduced by Daniel Bernoulli, is a logarithmic index intended to capture the psychological experience of satisfying a desire. Responding to a problem posed by his cousin, Nicholas, concerning how to price a sequential gamble with an unbounded expected value, D. Bernoulli's solution to the St Petersburg game appeals to the commonplace notion that, as your wealth increases, the pleasure you garner from an additional unit of reward decreases, all things considered. As a consequence, the disutility of a unit loss in wealth will be greater than the corresponding utility of a unit gain [1]—a property central to prospect theory [15] and variants [17, 24] introduced two centuries later.

In its original form, the *Bernoullian utility* of a prospective reward $x$ is

$$u_B(x) := \ln\left(\frac{w+x}{w}\right) \times c \tag{3}$$

where $c$ is a constant, $w$ an individual's current wealth, and $x$ a positive or negative reward. The natural log of the proportional change to one's wealth by receiving $x$, Bernoulli argued, *should* be inversely proportional to one's current wealth [1]. That said, few view non-linear utility as suited to guiding rational decision making. Instead, many continue to view deviations from the normative standard of linear utility as a regrettable but necessary means to describe stable patterns of human behavior [34].

Satisfying desires however is not a linear affair. Regarding your tenth pint as desirable as your first is a sign of addiction, not rationality. Yet, while one might agree that desire satisfaction is not and ought not be linear, it is doubtful the index is or should always be logarithmic. Descriptive models have behavioral data to fit, and normative standards can depend on context: the diminishing arcs of pleasure from accruing dollars and pints are not the same, nor should they be among sober adults.

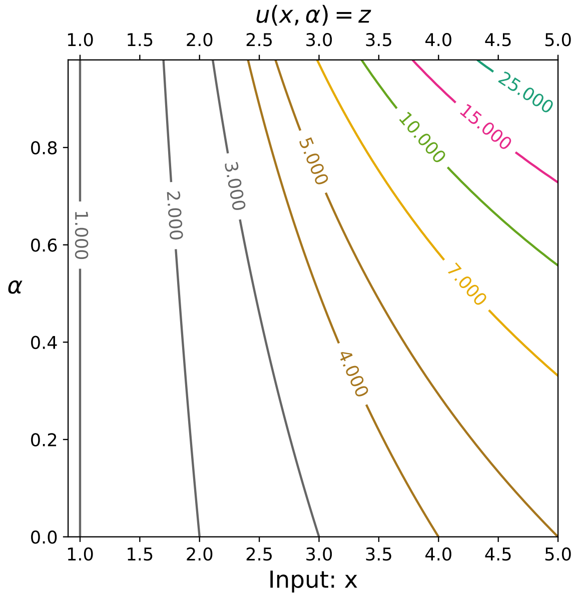


Figure 1: Discounted utility contour plot of nine discount curves for linear inputs $x$ from the set $\{1,2,3,4,5,7,10,15,25\}$, where the discounted curve for each $z$ is calculated by $u(x,\alpha)$ for all $0 \leq \alpha < 1$.

The challenge then is to construct a model for discounting desirable gambles that affords navigating the terrain of non-linear discounting with a model that is at once general yet simple enough to afford an interpretation of the degree of discounting that is applied. I propose a function of rewards with one adjustable parameter that is equivalent to linear utility (Equation 2), when no discounting is applied, and approaches natural log utility asymptotically as discounting approaches 1.

**Definition 1 (Discounted Utility)** *Let $\alpha \in [0,1)$ and $x > 0$. Define* discounted utility *as*

$$u(x,\alpha) := \frac{x^{1-\alpha} - \alpha}{1-\alpha}$$

Discounted utility takes a positive scalar reward $x$ and, depending on the value of the discounting parameter $\alpha$, discounts the desirability of $x$ to some degree. When no discounting is applied, $\alpha = 0$, then the utility of $x$ is linear, $u(x,\alpha) = x$. Alternatively, discounted utility approximates the natural logarithm of $x$ as $\alpha$ approaches 1. For completeness, one may stipulate that $u(x,1)$ is $\ln(x)$. Figure 1 illustrates the effects of the size of stake $x$ and size of discount rate $\alpha$ have on the discounted utility $z$, whereas Figure 2 illustrates the effective range of discounting for rewards $x$ between 1 and 100.

Every bounded gamble denotes a reward vector of positive or negative real numbers, $\mathbf{x}_f = (x_1, \ldots, x_m)$, but discounted utility is only defined for positive rewards. To accommodate the positivity condition for discounted utility,

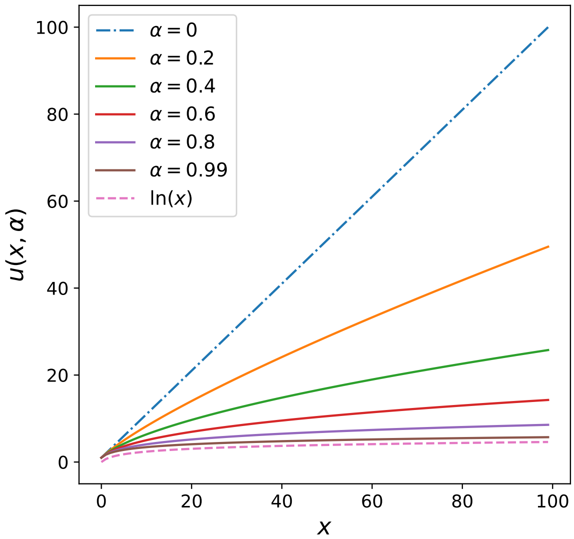


Figure 2: Discounted utility for input $x$, from 1 to 100, under six different discount values $\alpha$. When $\alpha = 0$, utility of $x$ is linear; as $\alpha$ approaches 1, utility approaches the natural log of $x$.

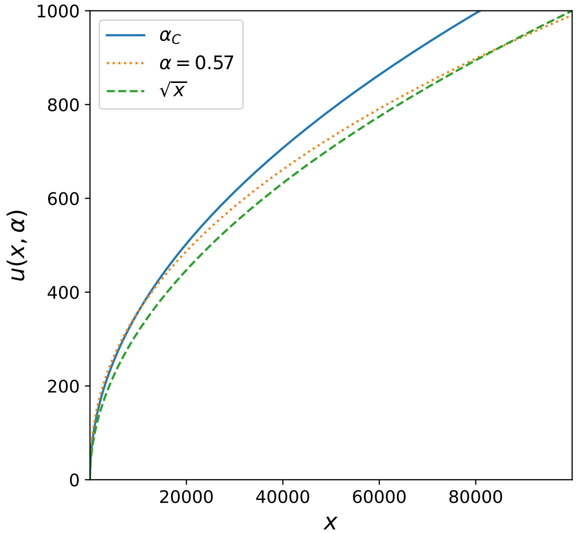


Figure 3: Discounted utility for input $x$, from 1 to 100,000, with Cramer's $\alpha_C$. Square root utility and discounted utility with $\alpha = 0.57$ are included for comparison.

assume that each gamble is adding positive or negative rewards to a current bank of positive wealth, $w$. Although one may relax this positivity condition and introduce a notion of debt, the basic model will instead treat 0 as an absorption point of insolvency and, for simplicity of exposition, assume a stake $w$ large enough to cover all possible losses of a one-shot gamble. Thus, a "solvent" transformation of any reward vector $\mathbf{x}_f$, $(w \pm x_1, \ldots, w \pm x_m)$, is everywhere positive, with $w$ rather than 0 denoting status quo ante. This positivity condition can be relaxed too with a policy for resolving insolvency.

It is so common now to ignore the issue of solvency that one may overlook the reason why this detail was set aside in the first place [8, §3.2.4]. Stake size does not matter for linear utility, by definition; so, designating 0 as status quo rather than an arbitrary $w$ is justified. But size matters to discounted utility. So, $w$ rather than 0 should be status quo.

Discounted utility is designed to take a linear, positive input $x$ and return a discounted valuation $z$ of $x$. Conversely, one may have a discounted $z$ and seek its undiscounted rate on a linear scale. This may be achieved by *reverse discounting*.

**Proposition 2 (Reverse Discounting)** *Let* $\alpha \in [0, 1)$ *and* $z > 1$. *The* inverse of discounted utility *is*

$$U'(z, \alpha) = (z - \alpha(z-1))^{-1/(\alpha-1)} \tag{4}$$

Discounted utility and reverse discounting are monotonic transformations of $x$ and $z$, respectively, thus are order preserving with respect to a given $\alpha \in [0, 1)$.

The family of discounted utilities can be extended by allowing $\alpha$ to be the value of a scalar function. For instance, Gabriel Cramer, a contemporary of Bernoulli's, proposed the square root transformation of rewards,

$$u_C(x) := \sqrt{x} \tag{5}$$

instead of a logarithmic transformation. One may extend discounted utility to accommodate an approximation of Cramer utility, balancing simplicity against faithfulness.

**Definition 3 (Almost Cramer utility)** *For any discounted utility u,* $x \geq 1$, *define* Cramer's alpha *by*

$$\alpha_C = \frac{\ln(4x)}{\ln(x^2)}$$

*Then,* $u(x, \alpha_C)$ *is called* almost Cramer utility.

The advantage of almost Cramer utility over an exact representation of square root utility is that Cramer's alpha is defined solely as a scalar function of $x$, dispensing with the imaginary unit $i$, $\pi$, and an integer constant an exact square-root $\alpha$ would require. Thus, almost Cramer utility is a simple but order preserving, dominating approximation of Cramer's square root utility (Figure 3).

## 4. Additive Discounting and Coherence

In the remainder I explore two strategies for discounting desirable gambles. The first approach, called *additive discounting* and described in this section, prioritizes the preservation of the mathematical properties of coherent sets of desirable gambles but dissociates the interpretation of desirability from the non-negative multiplicative scaling axiom, A3, from Section 2. Yet, because axiom A4 remains fixed, the additive structure of coherent sets of gambles is preserved. So, additive discounting maintains the mechanics of coherent inference by assuming that gambles only combine by a simple, additive dynamics.

The positive scale invariance axiom, A3, asserts that if a gamble is desirable, then so too is a stake of that gamble. Although the principal argument for viewing A3 and A4 as rationality postulates for desirability depends on expressing gambles within a linear utility scale [31, §2.3.4], positive scale invariance, on its own, is compatible with alternative notions of utility, include discounted utility.[3] Additive discounting instead uses discounted utility as a basis function to perform an element-wise transformation of the rays associated with each coherent gamble in $\mathbb{D}$.

A coherent set of gambles $\mathbb{D}$ is an object for inducing a coherent lower prevision, preserving partial preference, and performing coherent inference. Rather than regard an outcome vector $\boldsymbol{x}_f$ of each gamble $f$ in a coherent set $\mathbb{D}$ as an index of the desirability of $f$, the desirability of gambles is instead determined by a point-wise, additive transformation of elements of $\mathbb{D}$ by discounted utility. Alternatively, a collection of elicited gambles that are judged acceptable with respect to a discount rate $\alpha$ may be transformed to a linear scale by reverse discounting for the purposes of constructing a conic hull for coherent inference, if the elicited gambles are consistent and, once transformed, afford a natural extension. An exact transformation from elicited gambles to a convex cone requires an exact $\alpha$, but imprecision in $\alpha$ affords a form of sensitivity for the lower previsions induced by the collection of conic hulls produced from the admissible range of values for $\alpha$.

Figure 4 displays additive discounting of two solvent gambles, $x$ and $y$, for two different discount rates, $\alpha = 0$ (grey) and $\alpha = 0.5$ (color). Specifically, the $z$ axis represents the additive combination of a discounted $x$ and discounted $y$ by $u(x, \alpha) + u(y, \alpha)$. The grey hyperplane represents linear utility, for solvency adjusted values 1 to 10. The curved hyperplane is the additive combination of $x$ and $y$ with a discount rate of one-half.

**Proposition 4 (Consistency)** *Let $u(x, \alpha)$ be a map from $\mathbb{D}$ to $\mathbb{D}$ such that a discounted $\mathbb{D}^* \subseteq \mathbb{D}$. Then, a set of additively discounted $\mathbb{D}^*$ is consistent if $\mathbb{D}$ is coherent.*

The importance of preserving conic cones for coherent inference should be emphasized. A problem for Bernoullian utility is that it effectively depends on a distinguished currency scale. Since exchanging one currency into another is a linear transformation, but Bernoullian utility is logarithmic, switching one currency to another will involve exponentiating the currency exchange rate by a factor keyed to the original currency. Discounted utility is similarly without a base unit to underpin a conversion, which is why discounted utility *should not* be viewed simply as a type of currency conversion, and also why axioms A1 to A4 are supplemented by discounted utility transformation rather than altered.

3. I am grateful to an ISIPTA reviewer for pressing on this point.

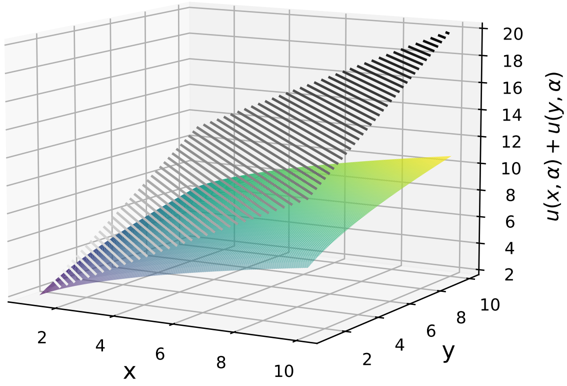


Figure 4: A linear utility hyperplane (grey) and discounted utility hypersurface (in color) for $\alpha = 0.5$, constructed from the additive combination of two reward vectors $x, y$, each with values spanning the real interval $[1, 10]$.

In short, the simple idea behind discounted utility is to maximize the utility of a reward rather than maximize the reward itself. Additive discounting extends the coherent sets of gambles framework by using discounted utility as a basis function to detach the operational meaning of desirability from the scale invariance axiom, A3, necessary for constructing conic hulls. By treating axioms A3 and A4 as closure conditions rather than rationality criteria per se, additive discounting moderates the role that linear utility plays in the theory of lower previsions.

## 5. Repeated Gambles

Gambles can be repeated, but repetition can mean different things. You might consider the outcome from playing a gamble $T$ times in a row to be the *sum* of the *expected* payoffs after each round of play, from 1 to $T$ rounds. Alternatively, you might view the result from playing the same gamble $T$ times to be the *product* of the *relative changes* in your wealth after each round of play. The first describes a *random additive process* for changes to your wealth over time and, as such, belongs to a class of stochastic processes where the central limit theorem figures in describing its long-term properties. The second describes a *random multiplicative process*, which is a stochastic process whose long-term behavior cannot be characterized in the same fashion because of the exponentially outsized effect of extremely rare events.

This section and the next address the dynamics under which gambles are combined. Random additive processes incorporate the additive structure that coherent sets of gambles inherits from expected utility theory [4]. The study of decision processes with imprecise transition probabilities

[37, 13, 30] and their limit behavior [7] are contributions to the study of additive dynamics. Random multiplicative processes depart from the additive combination condition (Axiom A4, Section 2), marking a more radical departure from the sets of desirable gambles framework.

Additive dynamics and multiplicative dynamics highlight a difference between the expectation value of a repeated gamble and the long-term time average of a single sequence of repeated gambles, and the conditions under which one average is equivalent to the other [18]. A sequence of gambles is *ergodic* if, given a sufficiently long time $t$, the observed empirical average of a gamble $f$ over that period of time, written $\mu(f(\omega),t)$, approaches a unique expectation value of $f$, $\mathbb{E}[f(\omega)]$, independent of initial conditions; that is, when

$$\begin{aligned}\lim_{t\to\infty}\mu\left[f(\omega,t)\right] &= \mathbb{E}\left[f(\omega)\right] \\ &= \mathrm{P}(f)\end{aligned} \tag{6}$$

where the reward of a gamble $f$ at $t$ is determined by the state $\omega\in\Omega$ realized at $t$ [2].

On the left-hand side of Equation 6, *Birkhof's Theorem*, the state $(\omega,t)$ is time-dependent, while the expectation $\mathbb{E}$ (linear prevision P) on the right-hand side relies on time-independent values associated with each $\omega\in\Omega$. When equation 6 holds, the time average is the same as the expectation value, so you can avoid integrating over time and instead integrate over the set of states, $\Omega$.

By emphasizing the gain or loss of a gamble *relative* to your current wealth in Section 3, I anticipated viewing the stakes of a gamble as a fixed proportion of current wealth at time $t$ rather than a fixed amount. To illustrate, consider a running example. Suppose you are offered a gamble $f$ on a coin toss that increases your current wealth $x(t)$ by 50% if the coin lands heads ($\omega_H$), and decreases your current wealth $x(t)$ by 40% if the coin lands tails ($\omega_T$). A time-dependent representation of this gamble is

$$f(\omega,t)=\begin{cases} f(\omega_H,t)=x(t)+0.5x(t) \\ f(\omega_T,t)=x(t)-0.4x(t)\end{cases} \tag{7}$$

For one round of play and an initial, status quo stake of €1 at $t_0$ for accepting the gamble, Equation 7 may be expressed as the gamble that results in your wealth increasing in absolute terms to €1.50 if heads or decreasing to €0.60 if tails. If the coin used to determine the states $\omega_H$ and $\omega_T$ is known to be fair and your current wealth at $t$ is $x(t)=1$, then the expected value of $f(\omega,t)$ is €1.05 and the variance is €0.20.

In a series of papers Ole Peters [21, 22], following an earlier observation by Maslov and Zhang [18, p. 382], argues that sequential decision problems are underspecified unless one explicitly addresses how the dynamics ought to be modeled. Asset price dynamics [14] and wealth dynamics more generally [11, 23] are non-stationary processes, as individual wealth tends to increase or decrease over time, not fluctuate around a stable mean. Yet, the presumption that Equation 6 holds is central to the Bayesian theory of inference and decision making with linear previsions, and this assumption largely carries over to the theory of lower previsions as well. Presuming that Equation 6 is satisfied without specifying how it is satisfied may be referred to as the *dogma of ergodicity*.

The procedure to model a dynamic process is to (*i*) convert a non-ergodic process into an "ergodic observable" of stationary, independent increments per unit of time; then (*ii*) maximize the expected value of these ergodic observables as an estimate for a time-average. In short, an ergodic observable is a growth rate, and the question becomes what (if any) transformation of a gamble $f$ is such that the expectation value of $f$, so transformed, is a suitable estimate of the time average of repeating $f$ according to that growth rate.

Following Peters and Gell-Mann [23], we can describe the process of converting non-ergodic additive and non-ergodic multiplicative repetitions of a gamble into ergodic observables for each dynamics. We assume throughout that the gamble $f$ is a stationary random variable.

**Additive Repetition.** Starting with an initial wealth of $x(t_0)$ at time $t_0$ before the first round of play, your accumulated wealth $x(t_0+T)$ after playing gamble $f$ for $T$ repetitions under a random additive process is calculated by

$$x(t_0+T)=x(t_0)+\sum_{t=1}^{T}f(\omega,t)$$

where $f(\omega,t)$ is the outcome from playing $f$ in the $t$'th round of play, for $T$ total rounds. The expectation value of your wealth $x(t)$ is not constant in time nor does the finite-time average converge to the expected value of $f$, 1.05. Thus, additive repetition is not ergodic.

Nevertheless, an ergodic observable exists in *additive changes* to wealth, $x(t+T)-x(t)$, which does not depend on $t$:

$$\lim_{T\to\infty}\frac{1}{T}\left[x(t+T)-x(t)\right]=\mathbb{E}_p\left[f(\omega)\right] \tag{8}$$

For simplicity, Equation 8 assumes the precise probability $p$ of heads and $1-p$ of tails is each one-half, but my general remarks about constructing an additive ergodic observable apply to lower previsions, too.[4] According to Equation 8, the expected outcome from playing $f$ repeatedly is that your wealth will grow exponentially in time by a factor of 1.05:

$$\mathbb{E}_p\left[x(t+1)\right]=\mathbb{E}_p\left[x(t)\cdot f(\omega,t+1)\right] \tag{9}$$

and the continuous growth rate is $\ln(1.05)$ or approximately 4.9%. It follows from Equation 9 that the expected

4. If the bias of the coin is unknown at $x(t)$, and you're considering the desirability of $f$, a consequence of coherence is that you ought to not accept $f$ if its lower prevision is less than 1, $\underline{\mathrm{P}}(x(t+1))<1$.

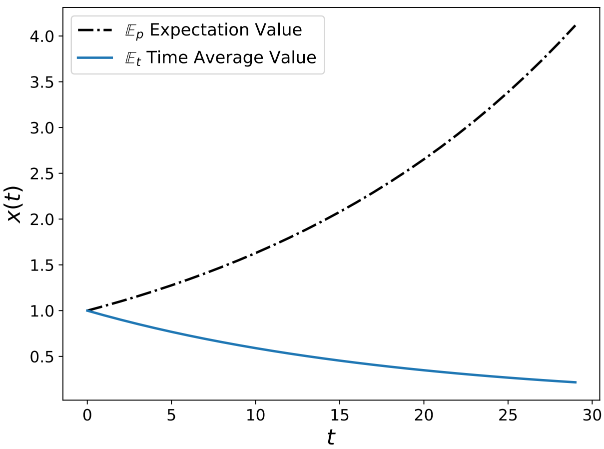


Figure 5: Plots of gamble $f$ described by Equation 7 with precise $p = 0.5$ and an initial stake $x(t_0) = 1$ repeated under two dynamics: (i) The expectation value calculated by Equation 8 yielding $\mathbb{E}_p \approx 4.12$ after $t = 30$ periods; (ii) The time average value calculated by Equation 11 returning $\mathbb{E}_t \approx 0.22$ after $t = 30$ periods.

value of your wealth after playing this gamble $T$ times is $x(t) \cdot \mathbb{E}_p [f(\omega,t)]^T$. ♢

So far I have taken the average of a binomial sequence, where $\mathbb{E}_p$ is the average of all possible outcomes computed by $(px_1 + (1-p)x_2)^T$, with $x_1, x_2 > 0$. The most probable $\hat{p}$ outcome of playing the gamble $f$ for $T$ rounds is a sequence that contains $Tp$ factors of $1.5x$ and $T(1-p)$ factors of $0.6x$, or $\hat{p} = (x_1^p x_2^{1-p})^T$. For additive processes, the average value of all possible outcomes, $\mathbb{E}_p$, is a good approximation of the most probable sequence $\hat{p}$, and converges as $T \to \infty$. For multiplicative processes, however, the ratio $\mathbb{E}_p / \hat{p}$ diverges exponentially in $T$ as $T \to \infty$ [25]. Figures 5 and 6 illustrate this divergence.

Continuing along the lines of [23], we turn next to the construction of a multiplicative ergodic observable.

**Multiplicative Repetition.** To describe a multiplicative process, first define per-round growth as

$$r(\omega,t) = \frac{x(t_0) + f(\omega,t)}{x(t_0)}$$

where $t_0$ is again the initial time before the first round of play. Since $f(\omega,t)$ is stationary and independent, so is $r(\omega,t)$. Then, a gamble is repeated multiplicatively if

$$x(t_0 + T) = x(t_0) \cdot \prod_{t=1}^{T} r(\omega,t),$$

which may be rewritten as

$$x(t_0 + T) = x(t_0) \cdot \exp\left[\sum_{t=1}^{T} \ln r(\omega,t)\right] \qquad (10)$$

Like additive repetition, neither is the expectation value of your wealth $x(t)$ constant in time nor does the finite-time average converge to the expected value of $f$ when repeated according to Equation 10. Unlike additive repetition, a *non-linear* transformation of $f$ is necessary to construct an ergodic observable of multiplicative changes in wealth over time. Specifically, under multiplicative dynamics, the ergodic observable is the rate of change in the logarithm of wealth.

$$\lim_{T\to\infty} \frac{1}{T} \ln\left(\frac{x(t+T)}{x(t)}\right) = \mathbb{E}_t[\ln f(\omega)] \qquad (11)$$

According to Equation 11, the expected outcome from playing $f$ repeatedly is that your wealth will decrease exponentially in time by a factor of approximately 0.95, as $x_1^p x_2^{(1-p)} = (1.5 \cdot 0.6)^{1/2} \approx 0.95 < 1$. Unlike additive ergodic observables, which are straightforward to construct for lower previsions, multiplicative ergodic observables are far less accommodating. Because their average is dominated by extreme events, the destabilizing role of extreme events is amplified if either the values of rewards $x_1$ and $x_2$ or the corresponding probabilities are far apart. Thus, the conservative behavior of lower and upper probabilities that licenses treating them as reasonable bounds on expectation values within additive structures and additive processes does not carry over to multiplicative processes. (See Figure 7 and Section 6.) ♢

Figure 5 illustrates the expected outcomes in wealth from playing the gamble $f$ thirty times with an initial stake of €1 and probability $p = \frac{1}{2}$ under additive ($\mathbb{E}_p$) and multiplicative ($\mathbb{E}_t$) dynamics. Assuming that wealth changes additively, the growth rate is the expectation value of $f$, $\mathbb{E}_p[f(\omega)] = 1.05$. Thus, $\mathbb{E}_p[f(\omega,30)]$ is expected to net a profit of approximately €3.22. Assuming that wealth changes multiplicatively, the time average growth rate is $\mathbb{E}_p[f(\omega)] = 0.95$. Thus, $\mathbb{E}_t[f(\omega,30)]$ is expected to yield a loss of approximately €0.78.

The choice of whether to play the gamble $f$ sequentially $T$ times is a choice between two different decision criteria: one that maximizes the ergodic growth rate assuming an additive dynamics; another that maximizes the ergodic growth rate assuming a multiplicative dynamics, which is modeled by changes in *logarithmic utility* (Equation 11). Peters and Gell-Mann refer to this second criterion as *Laplace's Criterion* [23], which I simplify here by assuming the scale for the rate of change is identical to the number of rounds of play.

**Definition 5 (Laplace's Criterion)** *Let $\gamma \geq 0$ be a fee paid in exchange for the gamble $f$ and $m = |\Omega|$. Then, to maximize the rate of change in logarithmic utility for a finite time period $T$ is to maximize*

$$\frac{1}{T}\sum_{i=1}^{m} p(\{\omega_i\})\left[\ln(x + f(\omega_i) - \gamma) - \ln(x)\right]$$

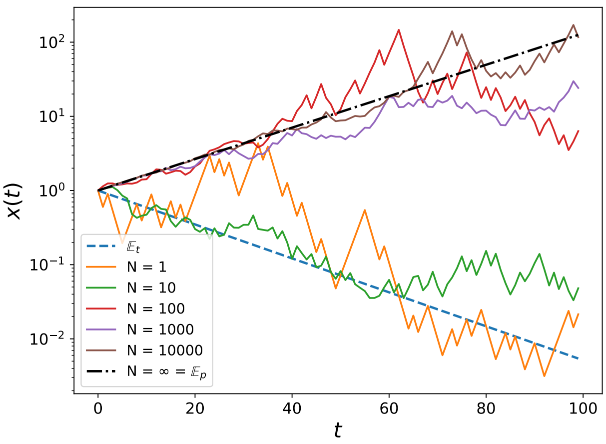


Figure 6: Simulations of portfolios of $N$ gambles (Eq. 7) played for $T = 100$ periods with $p = 0.5$ and resulting wealth $x(t)$ plotted on a log scale. The theoretical expectation values from Figure 5 are plotted as dashed lines. As the size $N$ of the portfolio increases the average of the portfolio approaches the expectation value $\mathbb{E}_p$.

Peters and Gell-Mann are critical of viewing multiplicative dynamics in terms of logarithmic utility. A gamble's mode of repetition is a fact about the gamble, not a psychological attitude about risk, even if the construction of an ergodic observable for multiplicative dynamics and Laplace's Criterion are formally equivalent. Whether you should expect to quadruple your stake or lose three-quarters of it rests on correctly modeling the dynamics. Rhetorical to [22] and fro [10] aside, the logarithm is capturing a critical feature of the multiplicative process, namely that the average of the $T$-fold product of $f$ is dominated by extreme "tail" events. The point to stress is that time-average performance is what you wish to approximate. Sometimes, such as for multiplicative repetition according to Equation 10, that approximation can be achieved by the expected value of an ergodic observable.

## 6. Adjustable Multiplicative Discounting

This said, maximizing the expected value and maximizing the expected log value are the two extremal values of discounted utility. However, to allay concerns that discounted utility is solely to do with individual preference, the discount rate $\alpha$ works similar to an informed prior in a dynamic setting: while you are free mathematically to choose any admissible $\alpha$ you like, the optimization objective to maximize your wealth compels you, rationally, to make an informed choice for the value of $\alpha$. And the information most relevant to you for rationally choosing $\alpha$ are facts about the gamble's dynamics, not your personal appetite for risk.

*Adjustable multiplicative discounting* departs from the additivity condition regulating how gambles combine over time, allowing for intermediate dynamics between pure additive repetition and pure multiplicative repetition. Unlike additive discounting, where the discount rate $\alpha$ was introduced without an operational interpretation, $\alpha$ does admit to rational calibration from knowledge about a gamble's dynamics.

Figure 6 plots the expected time-average of a sequence of gambles ($\mathbb{E}_t$) and the expectation value ($\mathbb{E}_p$), and simulated returns of portfolios consisting of $N$ gambles, all played for 100 rounds. As the portfolio size $N$ increases, the time average of the $N$-member ensemble time-average approaches the expectation value $\mathbb{E}_p$. In fact, the cross-over point where the ensemble size $N$ is large enough to be approximated by the expectation value for a finite period $T$ is on the order of $N \sim \exp(\eta T)$, for some constant $\eta$ [25]. In other words, the size $N$ of the portfolio needs to be on the order $e^T$ to effectively sample the extreme tail events over $T$ periods for the time average and ensemble average to align. But, even for just $T = 100$ rounds, constructing portfolios with $e^T$ many gambles is practically impossible.

In practice, real portfolios of multiplicative gambles will rarely if ever include $e^T$ gambles. Since the outcome of a portfolio will depend on the size $N$ of the portfolio and the number $T$ of iterations of play, the expected time average may be represented by discounted utility through simulations to estimate $\alpha$. This empirical method for fixing $\alpha$ through simulation is analogous to Bayesian predictive prior simulation, where you are choosing an $\alpha$ based on prior, perhaps partial, knowledge of the gamble dynamics and without foreknowledge of outcomes. As such, this empirical method for fixing $\alpha$ is normative, insofar as your decision to maximize your wealth $x(t)$ is influenced by the precision of your estimated payoff. These remarks suggest a generalization of Laplace's Criterion, which I call the *discounted utility criterion*.

**Definition 6 (Discounted Utility Criterion)** *Let $u$ be a discounted utility with $\alpha \in [0,1)$ fit by simulating portfolios of $N$ gambles played $T$ periods and $\gamma \geq 0$ a fee payed to play the gamble $f$ each period. Then, to maximize the rate of change in discounted utility for a finite time period $T$ is to maximize*

$$\frac{1}{T}\sum_{i=1}^{n} p(\{\omega_i\})(u([x + f(\omega_i) - \gamma], \alpha) - u(x, \alpha))$$

If there is uncertainty about the exact number $N$ of gambles in a portfolio, an upper and lower bound on $\alpha$ may be used to represent upper and lower expected values of a sequential gamble's payoff—with caveats. As a higher $\alpha$ corresponds to a smaller size $N$ ensemble, the ensemble time average is increasingly influenced by small fluctuations that are otherwise mean preserving as $N$ dwindles to 1. Nevertheless, one can put the haziness of lower-bound estimates of $\alpha$ into some context by comparing the order ef-

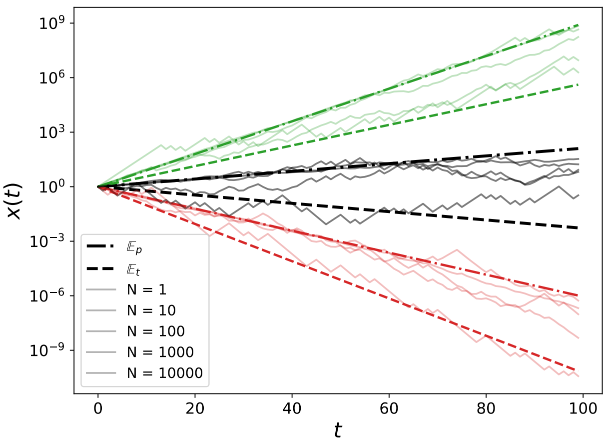


Figure 7: Simulations of portfolios of $N$ gambles (Eq. 7) played for 100 periods with $p = 0.3$ (red), $p = 0.5$ (grey), and $p = 0.7$ for the outcome $1.5x$, and resulting wealth $x(t)$ plotted on a log scale. The theoretical expectation values from Figure 5 are plotted as dashed lines.

fects of varying $\alpha$ to order effects of imprecise probability values.

Figure 7 compares the divergence in expected outcomes under additive and multiplicative dynamics from playing the gamble $f$ described in Equation 7 with probabilities $\boldsymbol{p} = (0.5, 0.5)$ (black) for the outcomes $(\omega_H, \omega_T)$, which are also displayed in Figure 6, to the same dynamics but with probabilities $\boldsymbol{p}^* = (0.7, 0.3)$ (green) and $\boldsymbol{p}_* = (0.3, 0.7)$ (red). Each of the three parameterized gambles includes (unlabeled) simulations for portfolios that includes $N$ = 1, 10, 100, 1000, and 10000 gambles for comparison.

If the probabilities $\boldsymbol{p}_*$ and $\boldsymbol{p}^*$ are viewed as witnessing lower(upper) and upper(lower) probabilities for $\omega_H(\omega_T)$, observe that the divergence in expected outcomes due to additive and multiplicative dynamics is swamped by the divergence in expected outcomes from each pair of upper and lower expectation values. Specifically, the difference between the expectation value and time average value after 100 trials is on the order of $10^4$, when $p = 0.5$, slightly less when $p = 0.7$ and slightly more when $p = 0.3$. However, the difference between the lower-bound time average when $\underline{\mathrm{P}}(\{\omega_H\}) = 0.3$ to the upper bound expectation average when $\overline{\mathrm{P}}(\{\omega_H\}) = 0.7$ is on the order of $10^{11}$. On this scale, discounted utility may provide a reasonable range of estimated outcomes for an uncertain ensemble size $N$ of gambles, even if the lower and upper estimates do not function as lower and upper bounds on expected outcomes.

## 7. Discussion

The desirable gambles framework is a proper extension of the credal sets approach to imprecise probabilities, and both are part of a tradition that seeks a unified, general account of coherent behavior. Even so, there are several alternative approaches to imprecise probabilities that do not fit neatly into the modern theory of lower previsions whose differences often spring from rejecting the behavioral, sure-loss avoidance and coherence foundations of the theory of lower previsions [16, 26, 33, 12]. This paper may be viewed as placing a foot in each of these traditions.

Additive discounting is a basis function approach to accommodating non-linear utility that aims to preserve the advantages of the theory of lower previsions while extending the theory's reach. I highlighted two issues with this approach. The first is that there is not a clear operational definition of the discount rate $\alpha$ that a complete normative theory might demand, although a non-trivial range of discount values may be easier to justify. The second is that discounted utility should not be viewed as a type of currency conversion, unless a case is made for a privileged scale for measuring desirability. Nevertheless, since discounted utility is an interpolation between linear and logarithmic utility, additive discounting enjoys an uncommonly high degree of interpretability for a linear basis function model.

In contrast to additive discounting, multiplicative discounting appears to upset the desirable gambles cart. The difference between multiplicative decision processes and additive decision processes involving imprecise transition probabilities is a difference of kind rather than of degree. Ensembles of multiplicative gambles on the order of $e^T$ can be reasonably estimated by expectation values. But, for $T$ of moderate length, the conditions for this equivalence are practically impossible to meet, so the chaotic properties of multiplicative random processes are unavoidable. Further, the discount rate $\alpha$ in discounted utility within adjustable multiplicative dynamics may be understood as a proxy for the size of an ensemble of gambles, and uncertainty bounds over the size of such a portfolio can afford a tighter interval estimation of outcomes than that offered by the theoretical expectation value and theoretical time average plotted in Figures 5, 6, and 7. Lastly, the interval estimates from moderate changes in probabilities, as one would find in a sensitivity analysis involving upper and lower probabilities, swamps the interval estimates induced by different dynamics. Put differently, for multiplicative decision processes, a lot rides on $p = 0.5$ meaning one-half rather than "I do not know".[5]

In closing, this preliminary study only scratches the surface of a number of interesting questions at the intersection of ergodicity breaking dynamical systems and imprecise probabilities. Although the proposal was advanced by numerical methods and a simple gamble with monetary wealth, applications to stochastic growth processes in biology, epidemiology, and the social sciences are conceivable, and analytical results within reach.

5. Compare to [20], which concerns the quality of one's information.

## Acknowledgments

My thanks to Lucas Böttcher, Jan Nagler and four anonymous ISIPTA reviewers for their constructive comments.

## References

[1] Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36, 1954. Trans. Louise Sommer, from "Specimen Theoriae Novae de Mensura Sortis", *Commentarii Academiae Scientiarium Imperialis Petropolitanae*, Tomus V, 1738, pp. 175–192.

[2] George D. Birkhoff. Proof of the ergodic theorem. *Proceedings of the National Academy of Sciences*, 17 (12):656–660, 1931.

[3] John C. Cox and Stephan A. Ross. The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3(1–2):145–166, 1976.

[4] Jasper De Bock and Gert de Cooman. Imprecise Bernoulli processes. In *International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU 2012)*, Advances in Computational Intelligence, pages 400–409, Berlin, Heidelberg, 2012. Springer.

[5] Gert de Cooman and Enrique Miranda. Weak and strong laws of larger numbers for coherent lower previsions. *Journal of Statistical Planning and Inference*, 138(8):2409–2432, 2008.

[6] Gert de Cooman and Erik Quaeghebeur. Exchangeability and sets of desirable gambles. *International Journal of Approximate Reasoning*, 53(3):363–395, 2012.

[7] Gert de Cooman, Filip Hermans, and Erik Quaeghebeur. Imprecise Markov chains and their limit behavior. *Probability in the Engineering and Information sciences*, 23(4):597–635, 2009.

[8] Bruno de Finetti. *Theory of Probability: A Critical Introductory Treatment*, volume I and II. Wiley, West Sussex, 2017 edition, 1974.

[9] Bruno de Finetti and Leonard J. Savage. Sul modo di scegliere le probabilità iniziali. *Biblioteca del Metron, Serie C*, 1:81–154, 1962.

[10] Jason N. Doctor, Peter Wakker, and Tong V. Want. Economists' views on the ergodicity problem. *Nature Physics*, 16(1168), 2020.

[11] Adrian Dragulescu and Victor Yakovenko. Statistical mechanics of money. *The European Physical Journal*, 17:723–729, 2000.

[12] Rolf Haenni, Jan-Willem Romeijn, Gregory Wheeler, and Jon Williamson. *Probabilistic Logics and Probabilistic Networks*. Synthese Library. Springer, Dordrecht, 2011.

[13] David Harmanec. Generalizing Markov decision processes to imprecise probabilities. *Journal of Statistical Planning and Inference*, 105(1):199–213, 2002.

[14] Ulrich Horst and Jan Wenzelburger. On non-ergodic asset prices. *Economic Theory*, 34(207–234), 2008.

[15] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47:263–291, 1979.

[16] Henry E. Kyburg, Jr. *The Logical Foundations of Statistical Inference*. D. Reidel, Dordrecht, 1974.

[17] Graham Loomes and Robert Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92(368):805–824, 1982.

[18] Sergei Maslov and Yi-Cheng Zhang. Optimal investment strategy for risky assets. *International Journal of Theoretical and Applied Finance*, 1(3):377–387, 1998.

[19] Enrique Miranda and Gert de Cooman. Marginal extension in the theory of coherent lower previsions. *International Journal of Approximate Reasoning*, 46 (1):188–225, 2007.

[20] Serafín Moral. Discounting imprecise probabilities. In Eduardo Gil, Eva Gil, Juan Gil, and María Á ngeles Gil, editors, *The Mathematics of the Uncertain*, volume 142 of *Studies in Systems, Decision and Control*, pages 685–697. Springer, Cham, Switzerland, 2018.

[21] Ole Peters. The time resolution of the St Petersburg paradox. *Philosophical Transactions of the Royal Society. Series A*, 369(1956):4913–4931, 2011.

[22] Ole Peters. The ergodicity problem in economics. *Nature Physics*, 15:1216–1221, 2019.

[23] Ole Peters and Murray Gell-Mann. Evaluating gambles using dynamics. *Chaos*, 26(023103):1–9, 2016.

[24] John Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3:323–343, 1982.

[25] Sydney Redner. Random multiplicative processes: An elementary tutorial. *American Journal of Physics*, 58(3):267–273, 1990.

[26] Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, NJ, 1976.

[27] C. A. B. Smith. Consistency in statistical inference (with discussion). *Journal of the Royal Statistical Society*, 23:1–37, 1961.

[28] Matthias C. M. Troffaes. Conditional lower previsions for unbounded random quantities. In J. Lawry, E. Miranda, A. Bugarin, S. Li, M.A. Gil, P. Grezegorzewski, and O. Hryniewicz, editors, *Soft Methods in Probability for Integrated Uncertainty Modelling*, Advances in Soft Computing, pages 201–209. Springer, 2006.

[29] Matthias C. M. Troffaes and Gert de Cooman. *Lower Previsions*. Wiley and Sons, Chichester, West Sussex, 2014.

[30] Matthias C.M. Troffaes. Learning and optimal control of imprecise markov decision processes by dynamic programming using the imprecise dirichlet model. In *Soft Methodology and Random Information Systems*, volume 26 of *Advances in Soft Computing*, pages 141–148, Berlin, Heidelberg, 2004. Springer.

[31] Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.

[32] Peter Walley. Towards a unified theory of imprecise probability. *International Journal of Approximate Reasoning*, 24:125–148, 2000.

[33] Kurt Weichselberger. The theory of interval-probability as a unifying concept for uncertainty. *International Journal of Approximate Reasoning*, 24 (2-3):149–170, 2000.

[34] Gregory Wheeler. Bounded rationality. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2018 edition, 2018.

[35] Gregory Wheeler. A gentle approach to imprecise probability. In Thomas Augustin, Fabio Cozman, and Gregory Wheeler, editors, *Reflections on the Foundations of Statistics: Essays in Honor of Teddy Seidenfeld*, Theory and Decision Library A. Springer, 2021.

[36] Gregory Wheeler. Moving beyond sets of probabilities. *Statistical Science*, 36(2):201–204, 2021.

[37] Chelsea C. White and Hany K. Eldeib. Markov decision processes with imprecise transition probabilities. *Operations Research*, 42(4):739–749, 1994.

[38] Peter M. Williams. Notes on conditional previsions. School of Mathematical and Physical Sciences, University of Sussex, 1975. Republished in *International Journal of Approximate Reasoning*, 44(3): 366–383, 2007.