

Dynamically Consistent Non-Expected Utility Preferences with Tuned Risk Aversion*

Berend Roorda[†] Reinoud Joosten[†]

June 22, 2015

Abstract

We propose the notion of Tuned Risk Aversion as a normative interpretation of non-expected utility preferences. It refers to tuning patterns of risk (and ambiguity) aversion to the composition of a lottery (or act) at hand, assuming an overall ‘budget’ for accumulated risk aversion over its stages. This makes the aversion level applied to a part intrinsically depending on the whole, in a way that turns out to be in line with frequently observed deviations from the Sure-Thing Principle. Uniqueness of updates is derived from a non-recursive form of consistency, in a general axiomatic setting, that also guarantees dynamic choice consistency under appropriate assumptions. The Allais paradox is used as leading example. Ambiguity aversion is illustrated by application to the 50:51 Example.

*We are grateful to Peter Wakker for valuable comments on a preliminary paper that has been presented at the FUR 2014 conference in Rotterdam.

[†]Faculty of Behavioural, Management and Social Sciences, Department of Industrial Engineering and Business Information Systems, University of Twente, the Netherlands. Email corresponding author: b.roorda@utwente.nl.

Keywords: preference updating, dynamic consistency, Sure-Thing Principle, consequentialism, Allais paradox, non-recursiveness, risk aversion, ambiguity aversion

1 Introduction

There is an abundance of evidence that risk attitudes towards a compound lottery, or act, cannot be properly understood in terms of risk attitudes towards each of its sub-lotteries separately, contrary to the implications of the von Neumann-Morgenstern expected utility framework. Since the famous example of Allais, decades of empirical and theoretical research have identified systematic aspects in human decision making that clearly violate the Independence Axiom underlying expected utility. This has resulted in several explanations in terms of psychological factors as regret, framing, and subjective perception of small probabilities, culminating in a variety of well-established frameworks for modeling so-called non-expected utility preferences. We refer to Machina and Viscusi (2013) for a recent overview on this topic. Despite agreement on the observed facts, there is still some controversy whether these deviations should be interpreted as biases of the human mind, comparable to optical illusions, or biases in the adopted principle itself, rendering it as less compelling than it seems to be at first sight; see Heukelom (2015) for a historical account. Our approach supports the latter view, by pointing at a straightforward explanation in terms of *Tuned Risk Aversion (TRA)*, which is inspired by recent findings in research on non-recursive valuation in the context of nonlinear pricing and risk measures in finance, see Roorda and Schumacher (2013, 2015) (henceforth RS13 and RS15) and the references therein.

The idea of TRA is best explained in the setting of a two-stage lottery, with only three degrees of risk aversion considered per stage: low, medium, high. To all

the nine combinations possible we can associate an induced ‘overall’ degree of risk aversion, applying to the lottery as a whole. Let us assume that a moderately risk averse person considers the following three combinations as not too conservative (a) medium in both stages, (b) first low, then high, and (c) the other way around. Variants are possible, of course, but the point is that in general there are level curves of overall risk aversion consisting of different patterns of distributing it. Tuned Risk Aversion, in this simple context, amounts to applying all three possibilities in this *tuning set*, and then selecting the one with minimum outcome. For example, if there is only risk in the second stage, (b) will be chosen, while (c) is most effective if only the first stage is risky. In this way risk aversion is tuned to the compound lottery as a whole, by ‘spending’ it ‘economically’, where it hurts most.

Natural as it seems, this example immediately raises serious concerns about rationality of the preferences induced: how to define a meaningful updating principle, and how to cope then with the dynamic inconsistencies that will inevitably emerge when the Sure Thing Principle is violated?

For the update rule, we rely on the notion of *sequential consistency*, introduced in Roorda and Schumacher (2007) and central in RS13 and RS15. In short, it only requires that certainty equivalents (ceqs) today are in the range of conditional ones tomorrow. This fits the idea of TRA much better than the far more restrictive axiom of recursiveness. We show that nevertheless it is strong enough to induce unique updates of initial preferences, in a quite general axiomatic framework for finite state acts that amounts to monotonicity, continuity, and a sensitivity condition on complete preference orderings (axioms A1-5). The implied update rule, which we call *fixed point updating*, turns out to be the one considered in Pires (2002). For TRA, with two stages, this rule amounts to applying the maximum tolerated level in the second stage. We also show that the existence of sequentially consistent updates imposes a restriction on initial preferences (axiom A6), which can be interpreted in

TRA as a ban on tuning risk aversion across mutually exclusive events.

Once updates have been defined, the issue of dynamic consistency emerges even sharper: how to rationalize the discrepancy between making conditional decisions in a certain state beforehand (according to a given initial preference ordering) and when that state actually materializes (according to its update)? The update may be mathematically well-defined, but does it make *sense*?

Our key observation at this point, leads to a subtle yet crucial modification of the standard definition of dynamic choice consistency (DCC), with stick-to-your-plan in a future state represented by a zero act, which is then compared to acts corresponding to exchanging the initial plan for an alternative. We show that this form of DCC is guaranteed under sequentially consistent updating for initial preferences that are super-additive (axiom A8), by deriving it from a static form of choice consistency that such preferences have (axiom A7). A more pragmatic condition, in terms of an externally given choice set, is proposed for other preference classes. Choice consistency is not addressed in RS13 and RS15.

From this perspective we take a final look at the interpretation of conditional ceqs, and their role in backward recursive evaluation. We argue that they generally do not function as replacement values of sub-acts, because it is inherent in most preference orderings that also other aspects of a sub-act matter. After all, this is precisely the point of the Allais paradox, which we use as a leading example in our exposition. These considerations lead to a revised notion of elementary building blocks of compound acts.

By an application to the the 50:51 example, introduced in Machina (2009), we illustrate the working of TRA in a context with both ambiguity and risk.

Finally we discuss related literature, emphasizing the contribution with respect to the multiple prior approach of Epstein and Schneider (2003), and the close connection with Pires' rule, also known as the axiom of conditional ceq consistency

in Eichberger et al. (2007) and Horie (2007). Some variants and extensions of our setting are formulated to enhance comparison with these and other frameworks, including an interpretation of updating in the Ellsberg paradox from our perspective.

This paper is organized as follows. The formal definition of TRA is given in Section 3, after the introduction and the description of the mathematical setting. Sections 4–6 form the axiomatic part, on resp. updating, dynamic choice, and recursion in ceqs. Each section ends with a leading example on the Allais paradox, and Section 7 contains the 50:51 example. Related literature and extensions are discussed in Section 8, and conclusions follow in Section 9. Proofs and a few technical results are collected in an appendix, which also contains a probability triangle on the Allais paradox as an additional illustration of the working of TRA.

2 Setting

We consider acts on a finite state space S , with consequences in $s \in S$ consisting of sub-acts with monetary outcomes in \mathbb{R} on a finite ‘final’ state space S_s . So acts take the form $f : \bar{S} \rightarrow \mathbb{R}$, with encompassing state space $\bar{S} := \cup_{s \in S} S_s$. The set of all acts is denoted as \mathcal{A} . The sub-act of an act $f \in \mathcal{A}$ in $s \in S$ is denoted as f_s , and \mathcal{A}_s denotes the set of all sub-acts in state s . When \bar{S} (S_s) is endowed with a probability measure, a (sub-)act is also called a (sub-)lottery. An act with the same consequence $c \in \mathbb{R}$ in all states in \bar{S} is called a *sure thing*, or the constant (act) c , or a sure amount c . The term generalizes to sub-acts in the obvious way, and we use the same notation for sure things on different state spaces. Contrary to common convention, acts with the same non-constant sub-act f_s in all $s \in S$ are *not* called constant. A *first stage act* f has no uncertainty after the first stage, i.e., all its sub-acts f_s are sure things. The set of all first stage acts on \mathcal{A} is denoted as \mathcal{A}^1 ; first stage acts are identified with mappings $S \rightarrow \mathbb{R}$.

The inequality $f \leq g$ means that $f(\bar{s}) \leq g(\bar{s})$ in all final states $\bar{s} \in \bar{S}$. Convergence of a sequence of acts is defined by identifying acts with vectors in \mathbb{R}^m with m the number of elements in \bar{S} . The numbers $\min f$ and $\max f$ denote resp. the minimum and maximum outcome of an act f , and $\text{range}(f) = [\min f, \max f]$.

We consider preference orderings \preceq on \mathcal{A} with the following standard properties. The symmetric and asymmetric part of \preceq are denoted by resp. \sim and \prec .

(A1) (*Weak order*) \preceq is complete and transitive.

(A2) (*Monotonicity*) If $f \leq g$, then $f \preceq g$.

(A3) (*Strict monotonicity for constants*) For $c, d \in \mathbb{R}$: $c < d$ implies $c \prec d$.

(A4) (*Continuity*) For a series of acts $(f_k)_{k \in \mathbb{N}}$ and act g , if $f_k \rightarrow f$, and $f_k \preceq (\succeq)g$ for all $k \in \mathbb{N}$, then $f \preceq (\succeq)g$.

Such orderings are called *regular*. It is well known that these are precisely the ones that can be represented by functions $\text{CE} : \mathcal{A} \rightarrow \mathbb{R}$ with the following properties:

(P1) (*Normalized on constants*) $\text{CE}(c) = c$.

(P2) (*Monotonicity*) If $f \leq g$, then $\text{CE}(f) \leq \text{CE}(g)$.

(P3) (*Continuity*) CE is continuous.

Functions satisfying P1 and P2 are called *certainty equivalence functions* (ceq functions). We call them *regular* if they also satisfy P3. Ceq functions on \mathcal{A}^1 are denoted by ce .

We use the symbol \preceq_1 for the state-dependent vector of preference orderings conditioned on the information after the first stage, i.e., $\preceq_1 = (\preceq_s)_{s \in S}$ with \preceq_s the ordering for sub-acts in state $s \in S$; $f \preceq_1 g$ means that $f_s \preceq_s g_s$ for all $s \in S$. The notion of regularity extends to \preceq_1 in the obvious way, and we use the symbol

cce for its ‘conditional’ ceq function. Concretely, $\text{cce} : \mathcal{A} \rightarrow \mathbb{R}^n$, with $\text{cce}(f)(s)$ the certainty equivalent of sub-act f_s under \preceq_s , and n the number of states in S ; $\text{cce}(f) \equiv c$ means $f \sim_1 c$. The outcome $\text{cce}(f)$ is identified with a first stage act in \mathcal{A}^1 , so that $\text{ce}(\text{cce}(f))$ is a meaningful expression.

3 Definition of TRA

TRA defines a preference ordering for compound acts in terms of an externally given family of preference functions per stage for a range of risk- and/or ambiguity aversion levels (*aversion levels* for short). We abstract from the way in which aversion is defined, and just assume a given parametrization of the degree of aversion with respect to a given ‘risk/ambiguity-neutral’ benchmark, having zero degree by convention, in the spirit of Yaari’s notion of *comparative risk aversion* (Yaari, 1969), see also Epstein (1999). More specifically, we assume that the following regular ceq functions are given:

- $\text{ce}_a : \mathcal{A}^1 \rightarrow \mathbb{R}$ for $a \in A$, with $0 \in A \subset \mathbb{R}$ denoting a range of aversion levels over the first stage, and ce_a non-increasing in a .
- $\text{cce}_b : \mathcal{A} \rightarrow \mathcal{A}^1 \simeq \mathbb{R}^n$ for $b \in B$ with $0 \in B \subset \mathbb{R}^n$ specifying given ranges of aversion levels for the second stage in each state, and cce_b non-increasing in each entry of b .

TRA is defined as taking the worst outcome over different patterns of aversion over the stages.

Definition 3.1 Tuned Risk Aversion, specified by a non-empty set $R \subset A \times B$, called the *tuning set*, corresponds to the preference ordering represented by the function $\text{CE}_R : \mathcal{A} \rightarrow \mathbb{R}$ given by

$$(3.1) \quad \text{CE}_R(f) = \inf_{(a,b) \in R} \text{ce}_a(\text{cce}_b(f)).$$

Some remarks are in order here. As the notation suggests, CE_R is indeed a ceq function, and CE_R is regular, i.e., satisfies the continuity property P3, under some additional regularity assumptions, see Lemma 10.1 in the appendix. In that case the infimum is achieved on compact tuning sets.

Every TRA-preference ordering CE_R has a maximum tuning set,

$$(3.2) \quad R^{\max} = \{(a, b) \in A \times B \mid ce_a(cce_b(f)) \geq CE_R(f) \text{ for all } f \in \mathcal{A}\}.$$

As we explained in the introduction, the crux of TRA is that the tuning set R need not be rectangular, but may reflect mutual restrictions between a and b .

At the outset, no restrictions are imposed on the benchmarks ce_0 and cce_0 , besides axioms A1-4. For lotteries, a standard choice is to let $a = 0$ and $b = 0$ correspond to taking (conditional) expectations under the given reference measure. More generally, one may choose ce_0 and cce_0 some linear functionals. A standard way to mirror positive to negative levels in nonlinear pricing is by setting $ce_{-a}(f) = -ce_a(-f)$, cf. Section 6.

Our definition of TRA can be extended in several obvious ways. Extension of TRA to more than two stages is addressed in Section 6, together with backward evaluation of CE_R . We have assumed a single aversion parameter for each stage, resp. $a \in A, b \in B$, but we could also incorporate different types of aversion in each stage, leading to higher dimensional sets A and B ; in fact, the example in Section 7 suggests an extension in this direction, see also Section 8.1. This indicates that a full axiomatic characterization of TRA may be complicated, and probably not particularly illuminating. In the discussion of consistency issues, however, we follow a purely axiomatic approach that does not refer to TRA, but takes starting point in axioms A1-4 only.

Example 3.2 We apply TRA to the lotteries of the Allais paradox, depicted in compound form in Figure 1. It has been well documented that many subjects prefer




Figure 1: The four lotteries related to the Allais paradox (payoffs in \$million).

a_1 over a_2 , and a_3 over a_4 , contrary to the Certainty-Independence Axiom that states that preferences should not switch if only the consequence c is changed from 0 to 1 while the sub-lottery is kept the same.

In our setting, the lotteries are viewed as acts $f : \bar{S} \rightarrow \mathbb{R}$ with final state space $\bar{S} = \{u, d, s'\}$, and also as compound acts on state space $S = \{s, s'\}$ with $s = \{u, d\}$.

We consider constant absolute risk aversion per stage, with a limit on the sum of their degrees,

$$(3.3) \quad CE_\gamma(f) = \min\{ce_a(cce_b(f)) \mid a, b \geq 0, a + b \leq \gamma\}.$$

Here ce_a is the certainty equivalent

$$(3.4) \quad -\frac{1}{a} \log(pe^{-ax_1} + (1-p)e^{-ax_2})$$

of a binary lottery with probability p on outcome x_1 and $1-p$ on x_2 under expected exponential utility $u(x) = 1 - e^{-ax}$; ce_0 is the expected value. The vector function cce_b is defined similarly in s (in s' it is the identity function). Notice that both stages are treated alike in CE_γ , so TRA does not rely on time dependency or extra model parameters.

The break-even point for the sub-lotteries, $cce_b(a_2)(s) = 1$, is reached for $b = 2.4$. So $a_1 \succ a_2$ if and only if $\gamma > 2.4$, e.g., $cce_b(a_2)(s) = 0.80$ for $b = 3$. It turns out that for the other pair of lotteries, risk aversion is most effective in the first stage, i.e.,

the minimum in (3.3) for a_3 and a_4 is achieved for $(a, b) = (\gamma, 0)$. Consequently, for the second stage expected values are considered, and hence a_3 is preferred over a_4 for all γ . So CE_γ is in line with the Allais preferences for $\gamma > 2.4$. Taking $\gamma = 3$, for example, yields $a_1 \sim 1$, $a_2 \sim 0.98$, $a_3 \sim 0.039$ and $a_4 \sim 0.036$.

This shows that the Allais preferences observed may be interpreted as an effect of tuning risk aversion. The example is continued in the next sections. As an additional result, the effect of TRA is depicted in a so-called probability triangle, in Section 10.2 of the appendix.

4 Sequential consistency and unique updating

In this section we describe the notion of dynamic consistency that we use, and its consequences for updating preferences. We follow an entirely axiomatic approach, independent of TRA. The results are applied to TRA at the end of the section.

Throughout this section \preceq denotes an (initial) preference ordering on \mathcal{A} , and $\preceq_1 = (\preceq_s)_{s \in S}$ a conditional one. In line with RS13, we impose the following relationship between \preceq and \preceq_1 . Recall that $f \preceq_1 g$ means that $f_s \preceq_s g_s$ for all $s \in S$.

Definition 4.1 (Sequential Consistency) We say that the pair \preceq, \preceq_1 is sequentially consistent, or that \preceq_1 is a sequentially consistent update of \preceq , if

$$(4.1) \quad c \preceq_1 f \preceq_1 d \Rightarrow c \preceq f \preceq d \quad (f \in \mathcal{A}, c, d \in \mathbb{R})$$

The following characterization further underlines the strong intuition of this axiom.

Lemma 4.2 *If \preceq_1 is regular, sequential consistency (4.1) is equivalent to*

$$(4.2) \quad f \sim_1 c \Rightarrow f \sim c$$

In terms of ceq functions CE for \preceq and cce for \preceq_1 , these criteria are respectively

$$(4.3) \quad \begin{aligned} & \text{CE}(f) \in \text{range}(\text{cce}(f)) \\ & \text{cce}(f) \equiv c \quad \Rightarrow \quad \text{CE}(f) = c. \end{aligned}$$

So sequential consistency requires that the ‘initial’ ceq of an act (under \preceq) must be in the range of the ‘sequential’ conditional ceqs (under \preceq_1), hence the name.¹ The characterization (4.2) expresses that if today one *foresees* indifference with a sure thing tomorrow, one *is* indifferent already today.

The theorem below shows that sequential consistency is strong enough to imply uniqueness of updates, under axioms A1-4 and a mild regularity condition on \preceq (axiom A5 below), and furthermore that the existence of a sequentially consistent update poses another, more substantial requirement on \preceq (axiom A6 below).

Let $f_s^c \in \mathcal{A}$ denote the act with sub-act f_s in state s of S , and sure thing $c \in \mathbb{R}$ in all other states of S .

(A5) (*c-Sensitivity*) If $f_s^c \sim c$, then $f_s^d \succ d$ for $d < c$ and $f_s^d \prec d$ for $d > c$ ($f \in \mathcal{A}$, $s \in S$, $c, d \in \text{range}(f_s)$).

(A6) (*c-Consistency*) If, for all $s \in S$, $f_s^c \sim c$ and $c \in \text{range}(f_s)$, then $f \sim c$ ($f \in \mathcal{A}$).

Axiom A5 compares a sub-act f_s in the context of different values for all other sub-acts being sure things of equal value in \mathbb{R} . The premise of the axiom, $f_s^c \sim c$, means that the sub-act f_s is like c under \preceq , in the context of c in other states; it is ‘neutral’ in the context of c . The axiom now imposes that \preceq recognizes that f_s rises above a worse context $d < c$. Similarly, for $d > c$, \preceq should sense that sub-act f_s is worse than d . The restriction of c, d to $\text{range}(f_s)$ in A5 is imposed to avoid that the worst- and best case preferences, $f \sim \min f$ and $f \sim \max f$, are excluded.

¹Sarin and Wakker (1998) uses the term sequential consistency for a recursiveness condition.

Axiom A6 is a consequence of sequential consistency, and hence implicitly motivated by that notion in the first place. The axiom expresses that if a set of sub-acts are each neutral in the context *separately*, then also *jointly*. In TRA, A6 can be motivated as a ban on tuning risk aversion between mutually exclusive events, as explained in the example below.

Theorem 4.3 *Let a preference ordering \preceq on \mathcal{A} be given that satisfies axioms A1-5. It has a unique update \preceq_1 determined by the fixed point update rule*

$$(4.4) \quad f_s \sim_s c \Leftrightarrow f_s^c \sim c \text{ with } c \in \text{range}(f_s) \quad (s \in S, f_s \in \mathcal{A}_s),$$

and this update is regular. If, in addition, \preceq satisfies axiom A6, this update is sequentially consistent, otherwise \preceq has no regular sequentially consistent update.

The fixed point update rule (4.4) determines c as the fixed point of the monotone mapping $c \mapsto \text{CE}(f_s^c)$, hence the name. It essentially the same as the rule in Pires (2002), see Section 8.2.

The following proposition captures the main intuition of the implications for TRA: updates must have the maximum tolerated aversion level in each state of S .

Proposition 4.4 *The pair $(\text{CE}_R, \text{cce}_\beta)$ is sequentially consistent if β is the maximum element of the set $R_1 := \{b \in B \mid (a, b) \in R \text{ for some } a \in A\}$, i.e., (i) $\beta \geq b$ for all $b \in R_1$ and (ii) $\beta \in R_1$.*

For the maximum tuning set (3.2) the given criterion is necessary as well, under some appropriate regularity conditions, see Section 10.6 in the appendix.

Example 4.5 The example CE_γ in (3.3) is defined as CE_{R_γ} with tuning set $R_\gamma = \{(a, b) \mid a, b \geq 0, a + b \leq \gamma\}$. The corresponding preference ordering \preceq^γ has sequentially consistent update \preceq_1^γ represented by cce_γ , which corresponds to the maximum tolerated level of risk aversion in the sub-lottery.

For an example without a sequentially consistent update, consider lotteries as in the Allais example, but with some binary sub-lottery also in state s' , and let $b(s)$ and $b(s')$ denote the aversion levels in resp. s and s' . The tuning set $R' = \{(a, b) \mid a, b \geq 0, a + b(s) + b(s') \leq \gamma\}$ does not satisfy axiom A6. Notice that R' involves the addition of aversion levels spent in mutually exclusive events. A6 forbids such a meaningless tuning of aversion levels.

Summarizing, CE_γ in (3.3) accommodates the Allais preferences for $\gamma > 2.4$, and has a unique update cce_γ that satisfies sequential consistency (4.1). The most delicate issue, that of dynamic choice consistency, is addressed in the next section.

5 Dynamic choice consistency

Although sequential consistency provides some basic form of dynamic consistency, there is still a tough nut to crack related to dynamic choices. Like any weakening of the Sure Thing Principle, it inherently leaves room for so-called dynamic inconsistencies: pairs of acts f, g can be found with $f \preceq_1 g$ yet $f \succ g$, as in the Allais paradox. That the preference of f over g is reversed for sure after the first stage seems hard to combine with a normative claim of a model.

The anomaly is perhaps felt most strongly in the following setting of a dynamic choice problem. Suppose f and g only differ in one sub-act, in state $s \in S$ say, with $f \succ g$ yet $f_s \prec_s g_s$, and one is offered the choice between f and g , with the option to switch after the first stage if s obtains. So one would prefer f over g precisely because of the difference in case s obtains, but actually then considers f worse!

We will not rationalize violations of the ‘stick-to-your-plan’ principle. Our point, however, is, that this principle may not be properly reflected by the requirement

$$(5.1) \quad f \succ g \quad \Rightarrow \quad f_s \succ_s g_s \text{ for some } s \in S \quad (f, g \in \mathcal{A}).$$

Even though we may prefer to obtain g_s rather than f_s in a state s , changing plans is

not about obtaining f_s again, but about abandoning the initial plan. If f is chosen initially, and then s obtains, the choice is in fact between doing nothing, or to replace it by the alternative. So what matters is the comparison of the sub-act $g_s - f_s$ with 0, rather than g_s with f_s .² We therefore propose the following alternative criterion for *Dynamic Choice Consistency* (DCC):

$$(5.2) \quad f \succ g \Rightarrow 0 \succ_s g_s - f_s \text{ for some } s \in S \quad (f, g \in \mathcal{A})$$

This argument does not rely on framing or endowment effects on preferences: we assume that the ordering \preceq_s itself is unaffected by committing to a plan. The chosen plan only sets the reference point in defining the to-be-compared acts, not subjectively, but *de facto*.

A similar line of reasoning motivates the following axiom for the initial ordering.

$$(A7) \text{ (Static Choice Consistency) (SCC)} \quad f \succ g \Rightarrow 0 \succ g - f \quad (f, g \in \mathcal{A}).$$

The criterion requires that if one prefers to obtain f rather than g , it cannot be that at the same time one has in mind that it would be attractive to exchange f for g . The following proposition states that choice consistency is preserved under sequentially consistent updating, and is guaranteed under the following axiom,

$$(A8) \text{ (Super-additivity)} \quad f \succeq c, g \succeq d \Rightarrow f + g \succeq c + d \quad (f, g \in \mathcal{A}, c, d \in \mathbb{R}).$$

A stronger result is obtained for the class of preferences satisfying

$$(A9) \text{ (c-Additivity)} \quad f \sim c \Leftrightarrow f - c \sim 0 \quad (f \in \mathcal{A}, c \in \mathbb{R}).$$

Proposition 5.1 *If \preceq satisfies A1-A6, and SCC (axiom A7), its sequentially consistent update \preceq_1 satisfies DCC (5.2). Under axioms A1-4, SCC is implied by A8, and if also A9 is satisfied, SCC is equivalent to A8.*

²In a setting with monetary outcomes, we consider $g_s - f_s$ as the most obvious way to represent the choice to return f_s for g_s . Some alternatives are discussed in Section 8.

For preference functions that do not satisfy SCC, weaker forms of choice consistency can be considered that restrict attention to an externally given choice set $\mathcal{C} \subset \mathcal{A}$. Suppose f^* is the unique optimal choice over \mathcal{C} , i.e., $f^* \in \mathcal{C}$ and $f^* \succ g$ for all other opportunities in \mathcal{C} . We say that \preceq exhibits *Static Plan Consistency (SPC)* with respect to \mathcal{C} if for the optimal ‘plan’ f^* it holds that

$$(5.3) \quad 0 \succ g - f^* \text{ for all } g \in \mathcal{C} \setminus \{f^*\}$$

If a choice set contains elements f, g that only differ in one sub-act, in state $s \in S$ say, this may be interpreted as a choice that can be postponed until the second stage, in case s obtains. For given $f \in \mathcal{C}$ and $s \in S$, we call $\mathcal{C}(f, s) := \{g \in \mathcal{C} \mid g_{s'} = f_{s'} \text{ for } s' \neq s\}$ the conditional choice set in s under ‘initial plan’ f ; for $f \notin \mathcal{C}$, it is the empty set. As criterion for *Dynamic Plan Consistency (DPC)*, for given choice set \mathcal{C} with unique optimal plan f^* , we take

$$(5.4) \quad 0 \succ_s g_s - f_s^* \text{ for all } g \in \mathcal{C}(f^*, s) \setminus \{f^*\} \quad (s \in S),$$

which is implied by SPC under sequentially consistent updating, in the same way as DCC by SCC.

We summarize the main implications for TRA in the following corollary.

Corollary 5.2 *If CE_R , defined by (3.1), is super-additive (A8), in particular when all single-stage ceq functions (ce_a for $a \in A$, cceb for $b \in B$) are super-additive, the represented ordering \preceq satisfies SCC (A7), and DCC is guaranteed under sequentially consistent updating. If CE_R exhibits SPC with respect to a choice set \mathcal{C} , then DPC w.r.t. \mathcal{C} follows under sequentially consistent updating.*

Example 5.3 We verify plan consistency under CE_γ with $\gamma = 3$ (see Example 3.2) for the choice set $\mathcal{C} = \{a_3, a_4\}$. SPC (5.3) requires that $h := a_4 - a_3 \prec 0$, and indeed $\text{CE}_\gamma(h) = -2.9$ (it turns out that the minimum in (3.3) is achieved for $a = 2.58$ and

$b = 0.42$ for h). DPC (5.4) is always implied by SPC for the sequentially consistent update \prec_s ; a direct verification yields $h_s \sim_s -3.97 = cce_b(h_s)(s)$ for $b = 3$.

Static and dynamic plan consistency also hold for choice set $\{a_1, a_2\}$, since $a_2 - a_1 = -h \prec 0$ and $-h_s \prec_s 0$ ($CE_\gamma(-h) = -0.02$, and $cce_3(-h_s)(s) = -0.20$).

This example does not satisfy the stronger condition SCC (axiom A7) of choice consistency with respect to any pair of acts. This follows from Proposition 5.1, and the fact that CE_γ satisfies axiom A9 but not A8. To see that it is quite possible to achieve SCC, we construct another example with super-additive Allais preferences.

Let ce'_a ($a \geq 0$) correspond to tolerating scaling of probabilities up to a factor $1 + a$,

$$(5.5) \quad ce'_a(g) = \min\{q_1x_1 + q_2x_2 \mid \frac{p_i}{1+a} \leq q_i \leq p_i(1+a)\},$$

with g the binary lottery with probability p_i on outcome x_i , $i = 1, 2$; cce'_b is defined similarly. Consider \preceq' corresponding to tuning set $R' = \{(a, b) \mid (1+a)(1+b) \leq 9\}$. As before, risk aversion is most efficient in the second stage for a_2 , and in the first stage for a_3, a_4 , so that the Allais preference ordering is maintained: $a_1 \sim' 1$, $a_2 \sim' 0.99$, $a_3 \sim' \frac{5}{90}$, $a_4 \sim' \frac{1.1}{90}$.

Since \preceq' is super-additive, SCC is guaranteed, and DCC as well for the sequentially consistent update \preceq'_1 represented by cce'_8 . So this example not only avoids dynamic choice inconsistency for the Allais lotteries, but for any choice set of acts on $\{u, d, s'\}$: whenever $f \succ' g$, then also $g - f \prec' 0$ and $g_s - f_s \prec'_s 0$ for some $s \in S$. No plan will be abandoned predictably under \preceq' .

6 Certainty equivalents and recursion

We obtained uniqueness of updating, and forms of dynamic plan consistency after a subtle modification of the standard definition. The question still remains how to interpret an update, if the corresponding conditional ceqs no longer provide the

replacement values of sub-acts that are required in a backward recursive evaluation. Moreover, if the initial ceq still can be interpreted as replacement value of entire acts, do we then treat future time instants on the same footing as current time?

We start with addressing a principal difference between ceqs and replacement values, independent of TRA. Consider an agent with preference ordering \preceq today, update \preceq_1 tomorrow, and a given act $f \in \mathcal{A}$ with $f \sim c$ and $f_s \sim_s c_s$ for $s \in S$. The replacement value d_s for f_s (under \preceq) is defined as the value for which $f \sim f'$ with f' the act f with f_s replaced by d_s ; it exists under our regularity assumptions, axioms A1-4, and we assume it is unique. The Sure Thing Principle demands that d_s only depends on f_s (consequentialism for replacement values), which would imply that $c_s = d_s$ under the fixed point update rule (4.4), so that recursiveness follows: c must be a function of c_s . We keep consequentialism, but only for ceqs: in the same way as c only depends on f , and not on acts foregone today, c_s only depends on f_s . The replacement value d_s , however, should not be identified with c_s , and it is inherent in \preceq that d_s may depend on the whole act f .

To clarify why, it is important to notice that preferences to obtain are different from preferences to offer. If $f \sim c$ means that ‘to the agent, obtaining f is indifferent to obtaining c ’, as we will assume now, this generally differs from the value $-c^* \sim -f$ he then will assign to *offering* the same act f .³ Similarly, c_s^* defined by $-c_s^* \sim_s -f_s$ need not be the same as c_s . So \preceq , the ordering *if it comes to obtaining*, induces another complete ordering \preceq^* for offering, defined by the reflection principle

$$(6.1) \quad f \preceq^* g \Leftrightarrow -g \preceq -f.$$

On the one hand, this ‘twin preference’ fully derives from \preceq , and inherits many of its

³We may assume that $c^* \geq c$, on the same grounds as for SCC (axiom A7). It may be illuminating to assume c -Additivity (axiom A9) and think of c and c^* as the agent’s bid- and ask price for f . In Section 8 we briefly indicate the connection with willingness to pay / accept, and describe some alternatives to the definition of c^* that avoid the use of $-f$.

features. In particular, each of the axioms A1-7 and A9 holds for \preceq if and only if it holds for \preceq^* , so that the reflection principle commutes with sequential updating. On the other hand, \preceq is concave if and only if \preceq^* is convex, so $\preceq^* \neq \preceq$ is the rule rather than the exception. Now the difference between c_s and d_s is that the one corresponds to obtaining tomorrow, the other to obtaining today. Considerations underlying the indifference that defines d_s , between obtaining f and f' today, naturally involve a comparison between *possessing* the sub-acts f_s or d_s already in case s will obtain, not only between obtaining them again tomorrow. So c_s is generally not the only aspect of f_s that determines d_s ; other aspects, such as c_s^* , may play a role as well. But then neither d_s need to be a function of c_s , let alone to coincide with c_s , nor c need to be a function of $(c_s)_{s \in S}$. Once this has been recognized, there is no compelling reason anymore to subject replacement values to consequentialism as a normative postulate.

For TRA the picture becomes more concrete. To evaluate $\text{CE}_R(f)$ backward recursively, it suffices to store the ‘extended act’ $f^B := (\text{cceb}(f))_{b \in B}$, with consequences in each state $s \in S$ consisting of ‘profiles’ $b(s) \mapsto \text{cceb}(f)(s)$. So under TRA, a compound act reduces to a first stage act with *profiles representing sub-acts*. On the profile in s , the ceq c_s of f_s corresponds to the maximum tolerated level for $b(s)$ in R , cf. Proposition 4.4, while the location of replacement value d_s on the profile depends on f . As illustrated below, it is located to the left of, or at, the point corresponding the level of risk aversion applied to evaluate CE_R : if $\text{CE}_R(f) = \text{ce}_{a^*}(\text{cceb}^*(f))$, then $\text{cceb}^*(f)(s) \leq d_s$.

To complete the recursion, and also to verify that initial time is treated at the same footing as future moments, *initial* profiles $f_0^\Gamma : \gamma' \mapsto \text{CE}_{\gamma'}(f)$ can be defined for $\gamma' \in \Gamma$, with Γ denoting a range of overall risk aversion levels, and $\text{CE}_\gamma = \text{CE}_R$ for some $\gamma \in \Gamma$. TRA then constitutes a backward recursion in terms of entire profiles, $f^B \mapsto f_0^\Gamma$, as illustrated below. This, eventually, resembles the ‘elementary’

recursion step in multi-stage TRA.

Example 6.1 In the example with CE_γ (3.3), the Allais lotteries can be reduced to single stage extended lotteries, with the sub-lottery in s represented by the profile

$$(6.2) \quad b \mapsto -\frac{1}{b} \log\left(\frac{10}{11}e^{-bx_1} + \frac{1}{11}e^{-bx_2}\right), \quad b \in [0, \gamma],$$

with x_1, x_2 the consequences in u, d , cf. (3.4); again for $b = 0$ the limiting value $\frac{10}{11}x_1 + \frac{1}{11}x_2$ is taken. The profile in state s' is, of course, the constant function equal to the consequence in that state. The function

$$\gamma' \mapsto \text{CE}_{\gamma'}(a_i), \quad \gamma' \in [0, \gamma]$$

can be interpreted as the current profile for act a_i in Figure 1, as it might have been considered in the past before the current state obtained. Negative aversion levels for b and γ' are not used in this example, but could be defined by the reflection principle (6.1).

This joint recursion in entire profiles cannot be reduced to point-wise recursion per aversion level. Replacement values range over the entire profile of the sub-lottery, in dependency of its context. For example, with $\gamma = 3$, as before, the same sub-lottery in a_2 and a_3 has replacement value 0.84 in a_2 and 4.55 in a_3 , corresponding to resp. $b = 2.86$ and $b = 0$ (the applied aversion levels b in TRA were resp. 3 and 0, cf. Example 3.2).

So different aspects are considered in different contexts. The Allais paradox has been *designed* to reveal that precisely this is a very natural thing to do. As we tried to argue in this leading example, it is not irrational either.




Figure 2: The four acts in the 50:51 example.

7 Tuning Risk and Ambiguity Aversion: the 50 : 51 example.

In Machina (2009) the so-called 50:51 example has been introduced, as an illustration of limitations in the Choquet Expected Utility (CEU) approach, introduced in Schmeidler (1989), to model the tradeoff between ambiguity and risk aversion. We take our starting point in the formulation of this example in Baillon et al. (2011), depicted in Figure 2.

CEU implies that f_1 is preferred to f_2 if and only if f_3 is preferred to f_4 . In Baillon et al. (2011) it is shown that the forward implication is induced in most other classes of ambiguity-averse preferences as well. The paradox is that the informational advantage of f_1 with respect to f_2 is much stronger than that of f_3 compared to f_4 , and hence it is natural allow for preferences \preceq with $f_1 \succ f_2$ and $f_3 \prec f_4$. We will show that TRA admits such a preference.

To keep things as simple as possible, we consider a preference function that is the minimum of two CEU functions. For ambiguity aversion, we consider ce_k for $k \in \mathbb{N}$ defined by

$$(7.1) \quad ce_k(g) = p^{k+1}x_1 + (1 - p^{k+1})x_2$$

with g the binary act with highest outcome x_1 with probability p and other outcome $x_2 \leq x_1$; cce_k is defined similarly. This is the so-called MINVAR($k + 1$) dis-

tortion measure, which has been proposed in Madan and Cherny (2010), together with several variants, as intuitive valuation measures in the context of bid-ask price modeling. The intuition is that the expected value is considered of the minimum outcome in $k + 1$ independent trials.

We take the uniform distribution as reference measure for the sub-acts, and choose $k = 1$. So the ambiguity penalty for the sub-acts amounts to a quarter of the spread of outcomes. There is no ambiguity in the first stage, so we take expected values of the outcome for both sub-acts, and define

$$(7.2) \quad U^{amb} = \frac{50}{101} \frac{E_1 + 3E_2}{4} + \frac{51}{101} \frac{E_3 + 3E_4}{4}.$$

The second CEU, reflecting risk aversion, is again obtained from exponential utility $u(x) = 1 - e^{-\beta x}$. The stepwise application of (3.4), with the same parameter in both periods, is equivalent to applying this utility once to four outcomes, so we take

$$(7.3) \quad U^r = -\frac{1}{\beta} \log\left(\frac{50}{101} \frac{e^{-\beta E_1} + e^{-\beta E_2}}{2} + \frac{51}{101} \frac{e^{-\beta E_3} + e^{-\beta E_4}}{2}\right).$$

These are both certainty equivalence functions, and we take their minimum outcome as final preference function V . It turns out that the values of f_1 and f_2 are equal for $\beta = 0.02365$. We choose $\beta = 0.015$ so that f_1 is preferred over f_2 . The corresponding values of the acts are given by

	f_1	f_2	f_3	f_4
U^{amb}	151	126.25	125.25	100.5
U^r	133.5	134.0	75.4	75.6
V	133.5	126.25	75.4	75.6

It follows that indeed $f_1 \succ f_2$ and $f_3 \prec f_4$, as desired.

V can be expressed as CE_R for some tuning set R , somewhat *ad hoc*, as follows. Define ce_0 and cce_0 corresponding to expected values, define cce_2 corresponding to $k = 2, \beta = 0$, cce_1 corresponding to $\beta = 0.015, k = 0$, ce_1 to $\beta = 0.015, k = 0$.

Then V corresponds to $R = \{(0, 2), (1, 1)\}$. We could add a third pattern $(2, 0)$ to increase the symmetry, to tolerate ambiguity aversion in the first stage, which then should be modeled so that it is without effect when probabilities are objective.

A more intuitive approach would involve two levels of aversion per period, one for risk, and one for ambiguity. A tuning set then consists of tolerated quadruples $(a, a'; b, b')$ with a, b degrees of risk aversion (as quantified by β in the example), and a', b' degrees of ambiguity aversion (as quantified by k).

We remark that TRA is certainly not the only way to cope with the 50:51 puzzle. We just showed that adding just one bit of TRA to CEU, nothing more, is enough, indicating that TRA extends scope of CEU in a relevant direction. We refer to Dillenberger and Segal (2015) for an alternative solution to this puzzle in terms of recursive preference functions, see also Section 8.3.

8 Related literature and extensions

In the literature on non-expected utility, a central role is played by so-called horse-roulette acts, having sub-acts modeled as lotteries with unambiguous (subjective or objective) probabilities, and a first stage with ambiguity. Therefore, we first relate our setup to some well-known approaches of this type that come close to TRA, and address a few modifications of our setup that facilitate the comparison. We then address the strong link with Pires' rule and Full Bayesian updating, and conclude this section with further discussion of related literature.

8.1 Comparison with MEU and recursive multiple priors

The class of Maxmin Expected Utility (MEU) preference functions has been introduced in Gilboa and Schmeidler (1989). In our setting, with outcomes in \mathbb{R} , they

take the form

$$(8.1) \quad \text{CE} = \text{ce}(\text{cce}(\cdot)) \text{ with } \text{cce}(f) = u^{-1}(E_1 u \circ f) \text{ and } \text{ce}(\cdot) = \min_{\mu \in \mathcal{Q}_0} u^{-1} E^\mu u \circ (\cdot)$$

with $u : \mathbb{R} \rightarrow \mathbb{R}$ a static utility function on outcomes, applied point-wise to outcomes of f in the expression $u \circ f$, E_1 an expectation operator conditional on S (under a second stage reference probability measure on S_s in each $s \in S$), and \mathcal{Q}_0 a collection of probability measures on S .⁴

Axioms A1-4 are satisfied (by ce , cce , CE) if u is strictly increasing and continuous. Axiom A5 is satisfied if for all $s \in S$, the set $\{\mu(s)\}_{\mu \in \mathcal{Q}_0}$ is bounded away from 0. In line with Epstein and Schneider (2003), also sets of probability measures can be considered in the second stage, so that $\text{cce}(\cdot)(s)$ is of the same form as ce in each $s \in S$, with sets \mathcal{Q}_s possibly depending on s . This results in the ‘multiple prior’ expression

$$(8.2) \quad \text{CE}(f) = \min\{u^{-1} E^{\mu'} u \circ f \mid \mu' \in \mathcal{Q}\},$$

with \mathcal{Q} satisfying two forms of rectangularity: (i) it is of the form $\mathcal{Q}_0 \times \mathcal{Q}_S$ and (ii) \mathcal{Q}_S is the rectangular product of $(\mathcal{Q}_s)_{s \in S}$. The extension provided by TRA is that the first form of rectangularity is relaxed: if \mathcal{Q}_S is the set of conditionals compatible with a risk-neutral μ_0 on S , then μ' on S that represent positive risk aversion may be combined only with strict subsets of \mathcal{Q}_S due to tuning restrictions. The fixed point update still corresponds to \mathcal{Q}_S , as in the rectangular case (Proposition 4.4), but the corresponding conditional ceqs are no longer replacement values of the sub-acts (Section 6). Axiom 6 requires that the second form of rectangularity is maintained,

⁴Formally, second stage lotteries are identified with probability distributions over a common outcome range X in each $s \in S$. Assuming $X \subset \mathbb{R}$, these distributions can be mapped, approximately, onto sub-acts in \mathcal{A}_s by taking $S_s = S'$ sufficiently large, and endowing S' with, e.g., the uniform distribution.

so that the fixed point update rule indeed produces a sequentially consistent update (Theorem 4.3).⁵

To obtain super-additive preferences (axiom A8) in (8.2), so that SCC is guaranteed (axiom A7), it is sufficient to choose a super-additive static utility function u . This essentially only leaves room for u of the form $u(x) = \lambda^+ x$ for $x \geq 0$, $u(x) = \lambda^- x$ for $x \leq 0$, with $\lambda^- \geq \lambda^+ > 0$. This is quite restrictive, but in fact $u(x) = x$ is not uncommon in a monetary context. By taking $\lambda^- > \lambda^+$ it still allows for a simple form of loss aversion, one of the distinctive features of (Cumulative) Prospect Theory (Tversky and Kahneman, 1992; Wakker and Tversky, 1993).

Notice that the utility function u in the above does not interact with the dynamics. To simplify formulas, we may specify acts in utility units (utils), and define $f' = u \circ f$, i.e., $f'(\bar{s}) = u(f(\bar{s}))$ for all $\bar{s} \in \bar{S}$, and represent \preceq on \mathcal{A} as \preceq' on the set \mathcal{A}' of all acts specified in utils, cf. also Gumen and Savochkin (2013). This is purely a matter of representation, but it may be noted that it does not make a difference whether axioms A1-6 are applied to \preceq , or directly to \preceq' . In TRA, tuning of risk aversion can then be specified on \mathcal{A}' , where possible aversion in the transformation of units by u is isolated from tuning aversion over stages.

Pursuing this idea somewhat further, some aspects of our framework may be generalized, as follows. Outcomes are often restricted to finite intervals, for instance in case the domain of u is bounded. Then expressions like $-f$, $g - f$ cannot be used, and our definitions of DCC, SCC, and \preceq^* require adjustment. From our viewpoint, stick to your plan still *has* to consider the opposite of obtaining an act, somehow. Instead of comparing $g_s - f_s$ with 0, as in our definition of DCC, we may consider

⁵Maccheroni et al. (2006) characterizes recursiveness for the broader class of Variational Preferences. The analogous result in the risk measure literature is Föllmer and Penner (2006, Thm. 4.5), on recursive convex risk measures. In RS15 sequential consistency is characterized for this class.

the choice between 0 and $u \circ g_s - u \circ f_s$, which is $g'_s - f'_s$ in the notation above. This leads to alternative definitions DCC' and SCC', in which the effect of replacement is measured in utils per state, rather than monetarily. These variants always hold for (8.2), as axiom A8 holds true in utils. Likewise, the reflection principle (6.1) for \preceq^* may be formulated in terms of transformed acts in \mathcal{A}' .

To conclude the subsection, we indicate how DCC and SCC may be adjusted without explicit reference to a static utility function u . The objects of choice are in fact pairs of acts $(f^{in}, f^{out}) \in \mathcal{A}^2$, meaning to obtain f^{in} and to return, deliver, or take a short position in f^{out} . Starting point is now a preference ordering \preceq^2 on such pairs. We saw two examples to obtain it from \preceq on \mathcal{A} : by identifying (f^{in}, f^{out}) with $f^{in} - f^{out}$, or with $u \circ f^{in} - u \circ f^{out}$, but there are many other possibilities. Following our line of reasoning for SCC, we may generalize this to the condition $(f, 0) \succ^2 (g, 0) \Rightarrow (g, f) \prec^2 (0, 0)$.

8.2 Related literature on updating

The fixed point update rule (4.4) is essentially the same as the notion of *conditional ceq consistency* in Eichberger et al. (2007), building on Pires (2002, Axiom 9), which is the forward implication in (4.4). In Pires (2002) the scope of the update rule is restricted to the Gilboa-Schmeidler framework. This includes the *dynamic* monotonicity axiom in (Gilboa and Schmeidler, 1989, Axiom 4), which is the condition (5.1) that we deliberately relax. Pires' rule closely relates to the Generalized Bayesian Rule in Walley (1991) and the Full Bayesian Updating Rule in Jaffray (1994). Our intended contribution is not so much an adjustment of these concepts, but an extension of their normative scope.

We have developed the rule without reference to additive or non-additive probabilities, only assuming axioms A1-5 and sequential consistency. This notion provides a further underpinning of the fixed point update rule, and leads to an additional re-

quirement (axiom A6) to ensure that the rule indeed produces a consistent update.⁶

A more detailed comparison leads to the following observations. We may assume the same space for sub-acts in each $s \in S$, i.e., that \mathcal{A}_s is independent of s , and assume that $\text{cceb}(s)$ is independent of s , in order to satisfy the axiom of state independence imposed in Pires (2002); Eichberger et al. (2007), following Gilboa and Schmeidler (1989). Obviously, \preceq_s satisfies the axiom of consequentialism.

We have restricted the attention to updates of the form \preceq_s , but the other setups also consider \preceq_E for subsets $E \subset S$, exactly as (4.4), but with s replaced by E . In particular \preceq_S is \preceq . This satisfies a compatibility property, called commutativity in Gilboa and Schmeidler (1989), which requires that \preceq_s can also be obtained as the update of \preceq_E with $s \in E$.⁷ The uniqueness and existence conditions in Theorem 4.3 generalize to \preceq_E in the obvious way.

For an extensive comparison of (4.4) with so-called f -Bayesian updating rules in (Gilboa and Schmeidler, 1993), we refer to Eichberger et al. (2010), in addition to the aforementioned references. This f -Bayesian update, for given $f \in \mathcal{A}$, is defined by $g \preceq_{E,f} h \Leftrightarrow g_E^f \preceq h_E^f$, whereas the rule (4.4), extended to events $E \subset S$, can be expressed as

$$g_E \preceq_E h_E \quad \Leftrightarrow \quad g_E^c \preceq c \preceq h_E^c \text{ for some } c \in \text{range}(g_E, h_E).$$

This may be viewed as a flexible form of f -Bayesian updating, in which the reference act f is taken a constant c that depends on the conditional acts that are compared.

⁶The notion of sequential consistency has been developed in an independent line of research on monetary valuations, which satisfy axiom A9 (they are also known as *monetary risk measures* under a different sign convention), see RS13; RS15. The refinement update described in these references corresponds to the fixed point update, but it is only applicable under A9, and no interpretation in terms of fixed points is given.

⁷Compatibility is addressed in (RS13, Prop. 4.6) and (RS15, Prop. 6.7) in more advanced settings.

The famous Ellsberg paradox has led to some alternative approaches to updating, in order to cope with dynamic consistency problems in the context of ambiguity, when the standard formalization (5.1) is adopted. The solution proposed in Hanany and Klibanoff (2007) is to adjust the update rule, by ruling out priors that cause a dynamic inconsistency. On the other hand, one may solve a violation of (5.1) by keeping the updated preference \preceq_s , and substitute the optimal conditional choice foreseen in s into f and g before comparing these acts at $t = 0$, according to the principle of ‘consistent planning’, as proposed in Siniscalchi (2009). We have argued, however, that under the new definition of dynamic consistency the anomalies may simply disappear. In particular, the preference in the example in Hanany and Klibanoff (2007), which is essentially the same as the example in Siniscalchi (2009, Section 2) satisfies all axioms A1-9 (it is of the form (8.2) with $u(x) = x$), and hence SCC and DCC are both guaranteed for any pairs of acts in the Ellsberg paradox, so that there is no need to adjust acts or update rules from our perspective.

8.3 Other topics in related literature

There are several other links with existing literature that we now briefly address.

Segal (1987) considers a recursive framework with two-stage lotteries, in which a probability measure obtains in the first stage that sets the odds for a set of outcomes in the second. As mentioned in Section 7, this has been successfully applied to the 50:51 example in Dillenberger and Segal (2015). Translated to our setting, each $s \in S$ is related to a probability measure Q_s on a finite set of outcomes x_1, \dots, x_k and sub-acts f_s are then the corresponding lottery under Q_s . Assuming a given probability measure Q on S , Segal’s approach is to replace each sub-lottery by its ceq under a given preference function V , and then evaluate the resulting first stage

lottery by (another, or possibly the same) preference function.⁸ A possible extension from the perspective of TRA is to allow for a range of preference functions V_b for the sub-lotteries, replace these in each s by their entire profiles $b \mapsto V_b(\cdot)$, and evaluate according to $\min_{(a,b) \in R} V'_a(V_b(\cdot))$ for a tuning set R .

An interesting result in this context has been obtained in Cerreia-Vioglio et al. (2015), providing an axiomatic characterization of so-called ‘cautious expected utility’, in a static setting with monetary prizes. This corresponds to considering the worst expected utility of lotteries over a set of utility functions u on final outcomes. Our dynamic framework offers a setting to incorporate this class of preferences in Segal’s approach, in the way just indicated.

Our discussion of twin preferences \preceq, \preceq^* was aimed at indicating limitations to the role of ceqs as replacement value. Obviously, they relate to the concepts of willingness to pay (WTP) and to accept (WTA), and to bid-ask price modelling. The vast literature on this subject involves many aspects that are far beyond our scope, such as a discussion of the Coase theorem, market efficiency and liquidity, elicitation methods, auctions. Incorporating insights from these fields may require the inclusion of wealth effects in our setting in the first place. A point of consideration arising from our analysis is that the frequently observed ‘WTP-WTA bias’ (see e.g. Machina and Viscusi (2013, Chapter 4)) may be rationalized to a larger extent than the name suggests. In this context it may be relevant to observe that the class of so-called α -MEU preferences corresponds to taking convex combinations $\text{CE}^\alpha := \alpha \text{CE} + (1 - \alpha) \text{CE}^*$, with CE^* representing \preceq^* in (6.1). CE^α satisfies axiom A1-6 if CE does, and for $\alpha = 1/2$, the ‘mid-price’ preference function is reflection symmetric:

⁸Reduction of compound lotteries, i.e., identification with the one-stage lottery obtained by first determining the initial probabilities on outcomes x_i is deliberately rejected. Our definition of compound acts as mappings $\bar{S} \rightarrow \mathbb{R}$ (here with $\bar{S} = S \times (S')^n$, $S' = S_s$ for all $s \in S$, consisting of k elements) respects this, because the staging structure is not lost in \bar{S} .

$(\text{CE}^{\frac{1}{2}})^* = \text{CE}^{\frac{1}{2}}$. In eliciting ceqs of subjects by choice lists it is important to realize whether one is after the ceqs corresponding to $\alpha = 0, \frac{1}{2}$ or 1. A subject that is fully consistent in adopting CE (from which all CE^α derive) will possibly give three different answers to three different questions. We also have argued that $\text{CE}^\alpha(f)$ may depend on profiles $(\text{c}ce^{\alpha'}(f))_{\alpha' \in [0,1]}$, rather than on its update $\text{c}ce^\alpha(f)$ alone.

The idea to keep track of more than one conditional utility level per state is also present in Vector Expected Utility (VEU) (Siniscalchi, 2009) and Expected Uncertain Utility (EUU) (Gul and Pesendorfer, 2014), but used differently.

Several of the aforementioned frameworks emphasize behavioral aspects of updating, in particular VEU and Prospect Theory, while we derived it straightforwardly from a consistency property, without reference to probabilities. However, our approach should not be seen as an interpretation that is alternative, or even opposite to the classical behavioral explanations in terms of framing, regret, endowment, and perception of small probabilities. On the contrary, it is very much in line with it. The different framing in the Allais lotteries ($c = 1$ vs. $c = 0$ in Figure 1) is precisely the reason for the different patterns of risk aversion between the first and second pair of lotteries, and the fact that a full loss is more painful in a_2 as compared to a_3 , nicely goes along with the fact that under TRA one is indeed more relaxed about that risk in the latter lottery. The gap between CE and CE^* may be given the interpretation of endowment, *inherent in* the preference ordering represented by CE. Tuning restrictions give room for strong emphasis on small (conditional) probabilities, as in the certainty effect, without letting it accumulate to excessive aversion over multiple stages. From this perspective, TRA may be seen as mechanism through which some aspects of these psychological factors have their effect. Our analysis seems to indicate that they have a more rational justification than generally believed, although we do not exclude that e.g., strong framing can lead to less efficient tuning of risk aversion. A proper treatment of behavioral aspects would

require a richer setting than we provided, in particular involving wealth effects and consumption. It is to be expected that further analysis of the empirical findings in this area will lead to refinements, extensions and adjustments of our framework.

9 Conclusions

We have introduced tuning of risk and/or ambiguity aversion (TRA) as a natural aspect of decision making under uncertainty, and described an axiomatic context in which it has a normative interpretation. Sequential consistency replaces recursiveness as a more flexible principle adjusted to TRA. It induces fixed point updating (also known as Pires' rule) as general updating principle that yields the unique candidate for a consistent update. This principle seems to capture the syntax of updating without reference to semantic information in probabilities. Instead, it only refers to the existence of a sure context in which the one sub-act stands out positively, the other negatively under a given initial preference. The key to further rationalization of TRA is an adjusted definition of dynamic consistency, resolving the conflict with consequentialism inherent in violations of the Sure-Thing Principle admitted under TRA. The final picture arising from our analysis is that of compound acts as composed of single-stage acts with consequences consisting of profiles rather than just one value. That this increased complexity need not come with extra model parameters, may enhance the formulation and testing of hypotheses in empirical research.

For future research, we would like to mention some important themes in non-expected utility that have not been addressed so far. Incomplete preferences can be defined, in the spirit of Aumann (1962) and Dubra et al. (2004), by combining a family of preferences in one overall preference, defined by the rule that an act f is 'overall'-preferred to g only when f is preferred to g under each preference in

the given family. An obvious idea in our setting is to consider incomplete orderings defined by the condition $CE_\gamma(f) \leq CE_\gamma(g)$ for a set of aversion levels γ . Inspiration for research in this direction may be derived from the results in Ok et al. (2012) on partial completeness.

Another possible application is prudence, see Eeckhoudt and Schlesinger (2006). Prudence seems inherent in TRA, in fact also if rectangular tuning sets are used for the two-stage lottery in the definition of prudence. Non-rectangular sets, however, may play a role in characterizing higher-order concepts related to the sign of the k -th derivative of utility functions for $k > 3$, such as temperance ($k = 4$) and edginess ($k=5$).

Finally, TRA may contribute to the analysis of non-recursive time preferences, in particular hyperbolic discounting (Phelps and Pollak (1968); Laibson (1997), see also Joosten (2014) for a recent application). One may first determine levels of risk aversion per stage that induce an appropriate discount rate for a reference set of acts. In TRA, tuning restrictions can then be imposed for lighter than recursive discounting over multiple steps. The results on compound risk measures in (RS13, Section 6), and the description of all convex risk measures with prescribed stepwise properties in (RS15, Section 7) may provide useful starting points for implementing this idea. More generally, sequential consistency and the induced fixed point update rule may provide a rational basis for giving long term consequences an appropriate weight in decision making, in a way that may be substantially different from the mechanically derived implications of local properties in a recursive approach.

10 Appendix

10.1 Lemma on continuity of CE_R

Lemma 10.1 CE_R defined in (3.1) is continuous on \mathcal{A} if

(r1) $A \times B$ is compact,

(r2) the mapping $(f, a) \mapsto ce_a(f)$ is continuous on $\mathcal{A}^1 \times A$, and

(r3) the mapping $(f, b) \mapsto cce_b(f)$ is continuous on $\mathcal{A} \times B$.

Then $CE_R(f) = CE_{\bar{R}}(f) = \min_{(a,b) \in \bar{R}} ce_a(cce_b(f))$, with \bar{R} the closure of R .

PROOF We start with the second claim. Conditions (r2-3) imply that

$$(10.1) \quad \text{the mapping } (a, b, f) \mapsto ce_a(cce_b(f)) \text{ is continuous.}$$

Due to (r1), R is bounded, and hence \bar{R} is compact. So $CE_R(f) = CE_{\bar{R}}(f)$ is the minimum of a continuous function over a compact domain, and by the Weierstrass Theorem it follows that the infimum is a minimum. The second claim of the lemma now follows.

For the first claim, we have to prove that if $f_n \rightarrow f$, then $CE_R(f_n) \rightarrow CE_R(f)$. In view of the claim just proved, we can write $y_n := CE_R(f_n) = ce_{a_n^*}(cce_{b_n^*}(f_n))$ and $y := CE_R(f) = ce_{a^*}(cce_{b^*}(f))$ for risk aversion levels in \bar{R} . Then $y_n \leq ce_{a^*}(cce_{b^*}(f_n)) =: z_n$, and hence $\limsup y_n \leq \lim z_n = y$, where the last equality follows from (10.1). It remains to show that $\liminf y_n \geq y$, or, equivalently, each limit point v of a converging subsequence $(y_n)_{n \in I \subset \mathbb{R}}$ satisfies $v \geq y$. Indeed, because \bar{R} is compact, such a subsequence must contain a sub-subsequence $(y_n)_{n \in J}$ with $J \subset I$, also converging to v of course, for which $(a_n, b_n)_{n \in J} \rightarrow (a', b') \in \bar{R}$, and hence $(a_n, b_n, f_n) \rightarrow (a', b', f)$. From (10.1) it then follows that $v = ce_{a'}(cce_{b'}(f)) \geq y$. \square

10.2 The Allais lotteries in a probability triangle

For the intuition, we depict the effect of TRA in a probability triangle, also known as a Marschak-Machina triangle, introduced in Marschak (1950), for the TRA preference (3.3) with $\gamma = 3$, see Figure 3. The fanning out effect in the lower region




Figure 3: The probability triangle with level curves of the TRA preference (3.3) with $\gamma = 3$ for the Allais lotteries. At the horizontal axis is the probability p_1 on outcome 0, on the vertical probability p_3 on outcome 5. The remaining probability $p_2 = 1 - p_1 - p_3$ is assigned to outcome 1. The lotteries of Figure 1 correspond to the following locations: $a_1 : (0, 0)$, $a_2 : (0.01, 0.1)$, $a_3 : (0.9, 0.1)$, $a_4 : (0.89, 0)$.

causes the outcome of preference orderings described in Example 3.2. We remark that the counter-intuitive North-West direction of curves in the upper region can be avoided if one minimizes over all three possible definitions of binary sublotteries in the Allais paradox.

10.3 Proof of Lemma 4.2

Necessity of the criterion (4.2) is obvious. To derive its sufficiency, consider f with $c \preceq_1 f \preceq_1 d$. Then there exist $g, h \geq 0$ so that $c \sim_1 f - g$ and $f + h \sim_1 d$ (continuity of \preceq_1 is essential here). By (4.2), then $c \sim f - g$ and $f + h \sim d$, and (4.1) follows from monotonicity of \preceq .

10.4 Proof of Thm. 4.3

We first prove that (4.4) defines a unique update \preceq_1 . Consider, for given $f_s \in \mathcal{A}_s$, the mapping $\rho : c \mapsto \text{CE}(f_s^c)$ on the domain $\text{range}(f_s) =: [l, r]$. By P1-3 for CE, cf. Section 2, ρ is continuous, $\rho(l) \geq l$ and $\rho(r) \leq r$. So ρ has a fixed point c' on this domain, i.e., there exists c' satisfying the right-hand side (rhs) of (4.4). Axiom A5 guarantees that such c' is unique, and hence that \preceq_s is uniquely determined by (4.4). This means that \preceq_1 is indeed unambiguously defined by (4.4).

Regularity of \preceq_s now directly follows from regularity of \preceq . In particular, \preceq_s is continuous, because for a series $f_{s,k} \rightarrow f_s$ in \mathcal{A}_s , with c_k the unique solution of the rhs of (4.4) for $f_{s,k}$, any converging subseries $(c_k)_{k \in \mathcal{I} \subset \mathbb{N}} \rightarrow c'$ yields $\text{CE}(f_s^{c'}) = c'$, by continuity of CE; so c' must be the unique solution of the rhs in (4.4), and hence the full series c_k is converging to c' .

Next we show that \preceq_1 defined by (4.4) is sequentially consistent if \preceq satisfies axiom A6, by deriving criterion (4.2). Let be given $f \in \mathcal{A}$ with $f \sim_1 c$. Then (4.4) implies that for all $s \in S$, $f_s^c \sim c$ with $c \in \text{range}(f_s)$, and by axiom A6 $f \sim c$, so that (4.2) follows.

It remains to show that if \preceq has a regular sequentially consistent update \preceq_1 , then \preceq must satisfy A6. Let an act $f \in \mathcal{A}$ be given with $f_s^c \sim c$ and $c \in \text{range}(f_s)$ for all $s \in S$. We have to prove that $f \sim c$. Consider an $s \in S$. As \preceq_1 is regular, there exists $c' \in \text{range}(f_s)$ such that $f_s \sim_s c'$, and hence $f_s^{c'} \sim_1 c'$. But then $f_s^{c'} \sim c'$ due to (4.2), while also $f_s^c \sim c$ by assumption, and axiom A5 implies that $c' = c$. Since $s \in S$ was arbitrary, $f_s \sim_s c$ for all $s \in S$, and, by (4.2), indeed $f \sim c$.

10.5 Proof of Prop. 4.4

We derive (4.3). Consider $f \in \mathcal{A}$ with $f \sim_s c$ for all $s \in S$, i.e., with $\text{c}e_\beta(f) \equiv c$. Condition (i) in the theorem implies that $\text{CE}_R(f) \geq \inf_{(a,b) \in R} \text{c}e_a(\text{c}e_\beta(f)) =$

$\inf_{(a,b) \in R} ce_a(c) = c$. From condition (ii) it follows that there exists $a \in A$ with $(a, \beta) \in R$, and hence $CE_R(f) \leq ce_a(cce_\beta(f)) = ce_a(c) = c$. So equality must hold, and (4.3) follows.

10.6 Extension of Prop. 4.4

In addition to the regularity conditions (r1-3) in Lemma 10.1 we consider the following sensitivity properties.

$$(r4) \quad ce_a(c_s^d) = d \Rightarrow c = d \quad (a \in A, c, d \in \mathbb{R}, s \in S).$$

$$(r5) \quad \text{There exist } c \in \mathbb{R}, f \in \mathcal{A} \text{ such that } cce_b(f) \equiv c \text{ and } cce_{b'}(f) \not\leq c \text{ for all } b' \not\leq b \\ (b \in B).$$

Lemma 10.2 *Under the regularity conditions (r1-5), CE_R has a regular sequentially consistent update if and only if R_1^{\max} has a maximal element β , namely cce_β .*

PROOF From Lemma 10.1 it follows that, under (r1-3), CE_R is continuous, that R^{\max} defined in (3.2) is a compact set, and that $CE_R = CE_{R^{\max}} = \min_{(a,b) \in R^{\max}} ce_a(cce_b(\cdot))$. Let a_{\max} denote the maximum value of a that occurs in R^{\max} , and let $\beta \in \mathbb{R}^n$ denote the least upper bound of R_1^{\max} (which is its maximum element if and only if $\beta \in R_1^{\max}$).

To see that (r4) now implies axiom A5, consider $f \in \mathcal{A}$ with $f_s^c \sim c$, i.e., $CE(f_s^c) = c = ce_{a^*}(cce_{b^*}(f_s^c))$ for some $(a^*, b^*) \in R^{\max}$. Condition (r4) implies that $cce_{b^*}(f)(s) = c$, that $cce_b(f)(s) \leq c$ for all $b \in R_1^{\max}$, and hence that $cce_\beta(f)(s) = c$. Axiom A5 now follows: for $d > c$, $CE(f_s^d) \leq ce_{a^*}(cce_{b^*}(f_s^d)) = ce_{a^*}(c_s^d) < d$, by (r4), and for $d < c$, $CE(f_s^d) \geq ce_{a_{\max}}(cce_\beta(f_s^d)) = ce_{a_{\max}}(c_s^d) > d$, again by (r4).

So we derived that under (r1-4), \preceq represented by CE_R satisfies axioms A1-5, and hence Theorem 4.3 is applicable. This implies that cce_β , which is indeed the outcome of the fixed point update rule (4.4), is the only candidate for a regular

sequentially consistent update, and it is such an update if and only if \preceq satisfies axiom A6. If $\beta \in R_1^{\max}$, axiom A6 holds, as we in fact already know from Proposition 4.4. It remains to deduce that $\beta \in R_1^{\max}$ from A6, or, more directly, from sequential consistency. From (r5) with $b = \beta$, it follows that there exist $c \in \mathbb{R}$, $f \in \mathcal{A}$ with $cce_\beta(f) \equiv c$ and $cce_{b'}(f) \not\geq c$ for all $b' \in R_1^{\max} \setminus \{\beta\}$. The sequential consistency criterion (4.3) implies that $CE_R(f) = c$, and hence $ce_{a^*}(cce_{b^*}(f)) = c$ for some $(a^*, b^*) \in R^{\max}$. From (r4) it follows that $cce_{b^*}(f) \equiv c$, and the inequality derived from (r5) implies that $b^* = \beta$, so indeed $\beta \in R_1^{\max}$. \square

10.7 Proof of Prop. 5.1

Under axiom A1-6, the sequentially consistent update \preceq_1 of \preceq exists. DCC is equivalent to the implication $g - f \succeq_1 0 \Rightarrow g \succeq f$. Indeed, $g - f \succeq_1 0 \Rightarrow g - f \succeq 0 \Rightarrow g \succeq f$, where the first implication follows from sequential consistency, and the second one from A7.

For the second claim, note that under axioms A1-4, \preceq is representable by a ceq function CE. To derive that A8 implies A7, consider $f, g \in \mathcal{A}$ with $f \succ g$. Then $CE(f) > CE(g) = CE(f + (g - f)) \geq CE(f) + CE(g - f)$ by axiom A8, so $CE(g - f) < 0$, and A7 follows (that continuity of CE is not needed here justifies the first claim of Corollary 5.2). For the last claim, first observe that axiom A7 is equivalent to the implication $CE(g - f) \geq 0 \Rightarrow CE(g) \geq CE(f)$. Under axiom A9, $g - f - CE(g - f) \sim 0$ for any pair $f, g \in \mathcal{A}$, and hence, by the implication just mentioned, $CE(g) \geq CE(f + CE(g - f))$. With $h := g - f$, it follows that $CE(f + h) \geq CE(f) + CE(h)$ for all $f, h \in \mathcal{A}$, which is axiom A8.

References

- R. J. Aumann. Utility theory without the completeness axiom. *Econometrica: Journal of the Econometric Society*, 30:445–462, 1962. [30]
- A. Baillon, O. L’Haridon, and L. Placido. Ambiguity models and the Machina paradoxes. *The American Economic Review*, 101(4):1547–1560, 2011. [20]
- S. Cerreia-Vioglio, D. Dillenberger, and P. Ortoleva. Cautious expected utility and the certainty effect. *Econometrica*, 83(2):693–728, 2015. [28]
- D. Dillenberger and U. Segal. Recursive ambiguity and machina’s examples. *International Economic Review*, 56(1):55–61, 2015. [22, 27]
- J. Dubra, F. Maccheroni, and E. A. Ok. Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115(1):118–133, 2004. [30]
- L. Eeckhoudt and H. Schlesinger. Putting risk in its proper place. *The American Economic Review*, pages 280–289, 2006. [31]
- J. Eichberger, S. Grant, and D. Kelsey. Updating Choquet beliefs. *Journal of Mathematical Economics*, 43(7):888–899, 2007. [5, 25, 26]
- J. Eichberger, S. Grant, and D. Kelsey. Comparing three ways to update Choquet beliefs. *Economics Letters*, 107(2):91–94, 2010. [26]
- L. G. Epstein. A definition of uncertainty aversion. *The Review of Economic Studies*, 66(3):579–608, 1999. [7]
- L. G. Epstein and M. Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003. [4, 23]
- H. Föllmer and I. Penner. Convex risk measures and the dynamics of their penalty functions. *Statistics and Decisions*, 24:61–96, 2006. [24]
- I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989. [22, 25, 26]
- I. Gilboa and D. Schmeidler. Updating ambiguous beliefs. *Journal of Economic Theory*, 59(1):33–49, 1993. [26]
- F. Gul and W. Pesendorfer. Expected uncertain utility theory. *Econometrica*, 82(1):1–39,

2014. [29]
- A. Gumen and A. Savochkin. Dynamically stable preferences. *Journal of Economic Theory*, 148(4):1487–1508, 2013. [24]
- E. Hanany and P. Klibanoff. Updating preferences with multiple priors. *Theoretical Economics*, 2(3):261–298, 2007. [27]
- F. Heukelom. A history of the Allais Paradox. *The British Journal for the History of Science*, 48(1):147–169, 2015. [2]
- M. Horie. Reexamination on updating Choquet beliefs. *KIER Discussion Paper*, 643, 2007. [5]
- J. Jaffray. Decision making with belief functions. In R. Yager, M. Fedrizzi, and J. Kacprzyk, editors, *Advances in the Dempster- Shafer Theory of Evidence*, pages 331–352. John Wiley, 1994. [25]
- R. Joosten. Social dilemmas, time preferences and technology adoption in a commons problem. *Journal of Bioeconomics*, 16(3):239–258, 2014. [31]
- D. Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478, 1997. [31]
- F. Maccheroni, M. Marinacci, and A. Rustichini. Dynamic variational preferences. *Journal of Economic Theory*, 128(1):4 – 44, 2006. [24]
- M. J. Machina. Risk, ambiguity, and the rank-dependence axioms. *The American Economic Review*, 99(1):385–392, 2009. [4]
- M. J. Machina and W. K. Viscusi. *Handbook of the Economics of Risk and Uncertainty*. Newnes, 2013. [2, 28]
- D. B. Madan and A. Cherny. Markets as a counterparty: an introduction to Conic Finance. *International Journal of Theoretical and Applied Finance*, 13(08):1149–1177, 2010. [21]
- J. Marschak. Rational behavior, uncertain prospects, and measurable utility. *Econometrica: Journal of the Econometric Society*, pages 111–141, 1950. [32]
- E. A. Ok, P. Ortoleva, and G. Riella. Incomplete preferences under uncertainty: Indecisiveness in beliefs versus tastes. *Econometrica*, 80(4):1791–1808, 2012. [31]
- E. S. Phelps and R. A. Pollak. On second-best national saving game-equilibrium growth.

- Review of Economic Studies*, 35(2), 1968. [31]
- C. P. Pires. A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2):137–152, 2002. [3, 12, 25, 26]
- B. Roorda and J. M. Schumacher. Time consistency conditions for acceptability measures, with an application to Tail Value at Risk. *Insurance: Mathematics and Economics*, 40: 209–230, 2007. [3]
- B. Roorda and J. M. Schumacher. Membership conditions for consistent families of monetary valuations. *Statistics & Risk Modeling*, 30:255–280, 2013. [2, 3, 4, 10, 26, 31]
- B. Roorda and J. M. Schumacher. Weakly time consistent convex risk measures and their dual representations. To appear in *Finance & Stochastics*, 2015. [2, 3, 4, 24, 26, 31]
- R. Sarin and P. P. Wakker. Dynamic choice and nonexpected utility. *Journal of Risk and Uncertainty*, 17(2):87–120, 1998. [11]
- D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989. [20]
- U. Segal. The Ellsberg paradox and risk aversion: An anticipated utility approach. *International Economic Review*, pages 175–202, 1987. [27]
- M. Siniscalchi. Vector expected utility and attitudes toward variation. *Econometrica*, 77(3), 2009. [27, 29]
- A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992. [24]
- P. Wakker and A. Tversky. An axiomatization of cumulative prospect theory. *Journal of risk and uncertainty*, 7(2):147–175, 1993. [24]
- P. Walley. *Statistical reasoning with imprecise probabilities*, volume 42. Peter Walley, 1991. [25]
- M. E. Yaari. Some remarks on measures of risk aversion and on their uses. *Journal of Economic Theory*, 1(3):315–329, 1969. [7]