

Consistent Preferences with Reversals

Berend Roorda* Reinoud Joosten*

This version: March 15 2017

Abstract

We propose an axiomatic framework for updating complete preference orderings without the monotonicity axiom. The induced preference reversals are justified on the basis of a normative interpretation that reconciles consequentialism with adequate forms of choice consistency. We conclude that preferences need not, and in some contexts even should not be recursive, thus confirming the intuition of the Allais and Ellsberg paradoxes.

Keywords: monotonicity, Sure Thing Principle, updating, dynamic consistency, choice consistency, non-expected utility, rationality, normative

JEL Classification: D81

1 Introduction

Monotonicity is at the root of normative frameworks for preference orderings on alternatives with uncertainty and risk. Abstracting from the differences among

*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.

mainstream approaches, such as the models of Savage and Anscombe-Aumann, it can be described as follows:

if A is preferred to B , conditioned on each event in an exhausting set of mutually exclusive events, then A is preferred to B . (a)

Here we avoid the more technical term ‘weakly-prefer’, which indicates that the preference need not be strict, but this is what is meant. Monotonicity has particularly strong consequences if the alternatives A and B are identical under all but one of the events:

if A and B only differ in event E , then A is preferred to B if and only if it is so conditioned on the event E . (b)

This has important implications in two directions. Forwardly, it provides an update rule, defining a conditional preference ordering for a given initial one. At the same time, it imposes a substantial restriction on the initial preference itself, to guarantee that the update rule is well-defined. If we adopt consequentialism, i.e., impose that the preference conditioned on event E only depends on the restriction of A and B to E , as we will do throughout the paper, we arrive at the ubiquitous Sure Thing Principle (STP) in Savage’s framework:

if A and B only differ in an event E , then A is preferred to B if and only if A_EC is preferred to B_EC for all C . (c)

Here A_EC, B_EC are the alternatives with the common part of A and B , outside E , replaced by some other common part, C . It may be noted here that the Independence Axiom, applicable to settings in the spirit of Anscombe-Aumann, implies the STP.

Normative frameworks that address updating commonly accept monotonicity as plausible, or even compelling. We take a different perspective, and do not consider

monotonicity as a *desideratum*. To motivate this deviating viewpoint, inspired by the findings in Roorda and Schumacher (2013, 2016) (henceforth RS13; RS16) is the main topic of our paper. In short, we argue that it is inherent in nonlinear preference orderings to depend on more aspects of conditional acts than their conditional value alone. We only adopt monotonicity in final outcomes as a normative axiom.

This allows us to weaken monotonicity to a far greater extent than generally regarded as appropriate in normative frameworks. In terms of the tripod of (a) a dynamic consistency principle, (b) an update rule, and (c) a soundness condition on initial preferences to make the update rule work, as sketched above for the monotonicity principle, the a-b-c of our proposal is

(a) Sequential Consistency: *if one is indifferent between A and a sure amount c , conditioned on each event in an exhausting set of mutually exclusive events, then one is indifferent between A and c .*

(b) Fixed Point Updating: *one is indifferent between A and a sure amount c conditioned on an event E if and only if one is indifferent between that amount c and A_{EC} .*

(c) the Equal Level Principle: *if one is indifferent between A_{EC} and a sure amount c , for all events E in an exhausting set of mutually exclusive events, then one is indifferent between that amount c and A .*

We have divided our exposition into three parts. In the first part, we develop a framework with existence and uniqueness of sequentially consistent updates, for a general class of complete preference orderings. We also show how this materializes for the concave and convex subclasses, in terms of their representations. For the sake of the argument, we strive for mathematical simplicity, and restrict attention to acts with monetary outcomes with finite range, on a finite outcome space. The

proposed update rule is in fact not new, but we believe that it is of independent interest to confirm it as a universal update principle that derives from sequential consistency. It substantially generalizes Bayesian updating, without any reference to probability or taste.

This, however, is only a stepping stone towards the main contribution we intend to make, in the second part: a rationalization of preference reversals, not only in the weak sense, as defined in Gilboa (2015), but in some respects also in the strong sense: that they *should* occur. Our argument relies on a careful distinction between the conditional value of a sub-act (consequentialist) and its replacement value (holistic), which in turn is justified by a nearly purely linguistic argument, that comparing acts is not the same as comparing obtaining acts. The proposed framework serves to make the ideas concrete, and to apply them to the Allais and Ellsberg paradoxes. Most ingredients in our exposition are by no means original, but their synthesis to a rationalization of reversals is, we believe, new. Furthermore, we obtain comforting dynamic consistency results that guarantee plan consistency of initially optimal plans, despite the lack of monotonicity, and hence recursiveness.

The third part is devoted to discussion of related literature and conclusions.

2 Setting and notation

2.1 Acts and class of preference orderings

We consider acts of the form $f : \Omega \rightarrow X$, with Ω a finite outcome space, and X a finite interval $[w, b] \subseteq \mathbb{R}$ of monetary outcomes. The set of all acts is denoted as \mathcal{A} , and \mathcal{A}^* is the set $-\mathcal{A}$, with outcome range $[-b, -w] =: X^*$. The case with $w = -b$, and hence $\mathcal{A} = \mathcal{A}^*$, will play a special role. The interval $[\min f, \max f]$ is denoted as $\text{range}(f)$. If an act f has $f(\omega) = c \in X$ on Ω , it is called a constant (act), and then we use the symbol c also for f . The (pointwise) mixture $\lambda f + (1 - \lambda)g$ is the

act with final outcome $\lambda f(\omega) + (1 - \lambda)g(\omega)$ in ω . An act is also called a lottery when an externally given probability measure on Ω is specified.

Our scope is the class \mathcal{P} of preference orderings satisfying the usual basic axioms.

Definition 2.1 \mathcal{P} is the class of preference orderings $\preceq \subset \mathcal{A} \times \mathcal{A}$ that satisfy

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

A2 (*Monotonicity in final outcomes*) If $f(\omega) \leq g(\omega)$ on Ω , then $f \preceq g$.

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

A4 (*Continuity*) For all $f \in \mathcal{A}$, the upper set $\{g \in \mathcal{A} \mid g \succeq f\}$ and the lower set $\{g \in \mathcal{A} \mid g \preceq f\}$ are closed.

Orderings in \mathcal{P} are called *regular*. The equivalence $f \sim c$ if and only if $V(f) = c$ defines a one-to-one correspondence between \mathcal{P} and the class of value functions $V : \mathcal{A} \rightarrow X$ that are continuous, monotone, and normalized, i.e., have $V(c) = c$. This V is called the (normalized) *value function* of \preceq , or the *certainty equivalence function* of \preceq . Proofs of these elementary facts are left to the reader. We will often abbreviate the term certainty equivalent to ceq.

2.2 State space and sub-acts

To streamline the exposition, updates are defined with respect to a state space S that corresponds to one degree of information under consideration. Formally, S is a partition of Ω . The state space S may be externally specified, or just hypothesized as a thought experiment in the decision making process.¹ How the case with several

¹As explained in Part III, for horse-roulette lotteries, S can be taken the end-states of the horse-lotteries (indeed commonly denoted as S), but also a partition of these end-states.

degrees of information can be reduced to one-shot updating with respect to a state space, is indicated by a reference to a compatibility result in Section 3.

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 . A (state) update of \preceq in s is a preference ordering on \mathcal{A}_s , commonly denoted as \preceq_s . For the vectors $(\mathcal{A}_s)_{s \in S}$ and $(\preceq_s)_{s \in S}$, we use the notation \mathcal{A}_1 and \preceq_1 , but $(f_s)_{s \in S}$ is simply identified with f . The (vector of) preference ordering(s) \preceq_1 is referred to as a (vector) update of \preceq . The definition of regularity extends to updates in the obvious way. We write $f_s h$ for the result of pasting sub-act f_s in state s in act h .

The axioms with reference to the state space S that we consider give rise to several subclasses of regular preference orderings in \mathcal{P} . In Part I we define the class \mathcal{S} with well-defined sequentially consistent updates, the concave class \mathcal{C} , and the class \mathcal{R} with dynamic risk aversion. In Part II we eventually define the classes \mathcal{M} and \mathcal{T} of resp. monetary and twin consistent preferences.

2.3 Notation for probability measures

Although our axiomatic setup does not refer to probability measures, except for the classes \mathcal{R} and \mathcal{M} , they are instrumental in the representation of concave preferences. Δ is the space of probability measures on Ω , $\Delta^+ = \{Q \in \Delta \mid Q(s) > 0 \text{ on } S\}$, and $\Delta_s^+ = \{Q \in \Delta \mid Q(s) > 0\}$. $E^Q f$ is the expected value of f under Q , Δ_s and E_s^Q are the corresponding conditional versions in state s , and $E_1 f$ is $(E_s f)_{s \in S}$. Further, TQ_s is the probability measure obtained by pasting Q conditioned on s in T , formally defined by $E^{TQ_s} f = E^T(1_s E_s^Q f + (1 - 1_s) E_1^T f)$. In the context of risk aversion, we use P for the reference measure.

All notation generalizes from states s to events $E \subseteq \Omega$ in the obvious way. The usual double role of the symbol E should not cause confusion.

Part I

Updating without the monotonicity axiom

We formulate an axiomatic framework for unambiguous updating, based on the notion of sequential consistency, and the corresponding Equal Level Principle as generalization of the Sure Thing Principle. The general result is made concrete for the concave and convex subclasses, in terms of their representations. The update formulas are simplified under additional assumptions on risk aversion.

3 Main axioms and update rule

The following axioms form the cornerstone of our framework. They apply to $f \in \mathcal{A}$.

S1 (*Sequential Consistency*) If $f_s \sim_s c$ on S , then $f \sim c$.

S2 (*Equal Level Principle*) If $f_s c \sim c \in \text{range}(f_s)$ on S , then $f \sim c$.

S3 (*c-Sensitivity*) If $f_s c \sim c$, then $f_s d \succ d$ for $d < c$ and $f_s d \prec d$ for $d > c$, for all $s \in S$ and $c, d \in \text{range } f_s$.

Definition 3.1 \mathcal{S} is the subclass of preferences in \mathcal{P} that satisfy axioms S2 and S3.

The first axiom, *Sequential Consistency*, is our key notion for consistent updating, replacing the common notion of monotonicity,

$$f \preceq_s g \text{ on } S \quad \Rightarrow \quad f \preceq g \quad (f, g \in \mathcal{A}), \quad (3.1)$$

The notion of sequential consistency has been developed and analyzed in a long-standing research line on risk measures and valuations, see Section 14. It is equiva-

lent, in \mathcal{P} , to the condition

$$c \preceq_1 f \preceq_1 d \quad \Rightarrow \quad c \preceq f \preceq d. \quad (3.2)$$

In other words, *values should be in the range of their sequential updates*. Axiom S2 is the corresponding weakening of the STP, deemed the *Equal Level Principle*, that characterizes existence of consistent updates, under the sensitivity condition of axiom S3 that guarantees their uniqueness.

As shown in the theorem below, these axioms lead to the following update mechanism, which we call *fixed point updating (fpu)*:

$$f_s \sim_s c \quad :\Leftrightarrow \quad f_s c \sim c \text{ with } c \in \text{range}(f_s). \quad (3.3)$$

We call \preceq_s a *fixed point update* of \preceq (in state s) if it satisfies the forward implication in (3.3); it satisfies (3.3) if and only if it is the unique one.

Theorem 1 *A preference ordering \preceq in \mathcal{P} has unique fixed point updates \preceq_s on S if and only if \preceq satisfies axiom S3, and then \preceq_s is given by (3.3), and regular. The (vector) update $(\preceq_s)_{s \in S} =: \preceq_1$ is then sequentially consistent (axiom S1) if and only if \preceq also satisfies axiom S2, otherwise \preceq has no regular sequentially consistent update.*

So \mathcal{S} is the subclass of \mathcal{P} for which the fixed point update (3.3) is well-defined and produces the sequentially consistent update.

It may be illuminating to compare the implications of axiom S2 and the STP for a strictly monotone preference ordering \preceq on acts with three outcomes, (x, y, z) , and two states (s, s') , corresponding to resp. the first two outcomes and the third. The STP requires that $(x, y, z) \sim (c, c, z)$ either for all z or none. Axiom S2 amounts to the implication that if $(x, y, c) \sim c$ and $(c, c, z) \sim c$, then $(x, y, z) \sim c$, which is void, since $z = c$ when $(c, c, z) \sim c$. The knowledgeable reader will immediately recognize what this means for the Allais paradox. Note that axiom S3 is satisfied,

for instance, when the induced value function $V(x, y, z)$ has both the third partial derivative strictly bounded by 1, as well as the sum of the first two.

We conclude this section by a remark on compatibility of updates. The notation and preceding results generalize from states $s \in S$ to events E in a partition of S in the obvious way. In particular, under the analogues of axiom S2 and S3, the consistent update \preceq_E is then determined by (3.3), with s replaced by E . 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 $s \in E$. Compatibility is addressed in (RS13, Prop. 4.6) and (RS16, Prop. 6.7) in technically more advanced settings.

4 The concave subclass and its representation

Concave preferences are those that satisfy

C1 (*Concavity*) If $f, g \succeq c$, then $\lambda f + (1 - \lambda)g \succeq c$ for all $\lambda \in [0, 1]$.

In other words, their upper contour sets,

$$\mathcal{A}^c := \{f \in \mathcal{A} \mid f \succeq c\}, \quad (4.1)$$

are convex, for all $c \in X$. Concavity is strongly related to risk- and ambiguity aversion, also when it is defined in terms of probability mixtures. The corresponding property for value functions is quasi-concavity:

$$V(\lambda f + (1 - \lambda)g) \geq \min\{V(f), V(g)\}.$$

In the representation lemma below, which is a simplified yet slightly different version of the results in Cerreia-Vioglio et al. (2011b), we use the following regularity condition.

C2 (*Sensitivity for constants*) $f + d \succ f$ for $d > 0$, $f + d \in \mathcal{A}$.

Lemma 4.1 *A preference ordering \preceq is concave (axiom C1) and regular if it can be represented as*

$$f \sim c \quad \Leftrightarrow \quad \min_{Q \in \Delta} E^Q f + \theta(c, Q) = c, \quad (4.2)$$

for some function $\theta : X \times \Delta \rightarrow \mathbb{R}^+$ that is (i) continuous in the first argument, (ii) satisfies $\min_{Q \in \Delta} \theta(c, Q) = 0$, for all $c \in X$, and (iii) has the property that $\theta(c, Q) - c$ is equi-strictly decreasing.² Under axiom C2, \preceq is concave and regular only if

$$\theta(c, Q) = \max\{c - E^Q f \mid V(f) = c\} \quad (4.3)$$

is such a representation, and this representation is minimal.

Concave preference orderings can also be represented by R -representations,

$$V(f) = \min_{Q \in \Delta} R(E^Q f, Q) \quad (4.4)$$

with $R : X \times \Delta \rightarrow X$. These can be obtained from θ -representations, and vice versa, by the equivalence

$$R(m, Q) = c \quad \Leftrightarrow \quad \theta(c, Q) = c - m. \quad (4.5)$$

It follows that the minimal R -representation of \preceq in \mathcal{C} is given by

$$R(m, Q) = \max\{V(f) \mid E^Q f = m\}. \quad (4.6)$$

The proof of Lemma 4.1 is in terms of R -representations, and gives the properties of R analogous to (i)-(iii) in the lemma.

5 The fixed point update for concave preferences

We now come to the translation of the fixed point update rule in terms of representations of concave preferences. In addition to axiom C2, we impose another regularity condition to guarantee the relevance of each state.

²This means that for all $\varepsilon > 0$, there exists $\delta > 0$ such that, for all $Q \in \Delta$, $|d - c| \geq \varepsilon$ implies $|\theta(d, Q) - d - (\theta(c, Q) - c)| \geq \delta$.

C3 (*Sensitivity per state*) $b_s c \succ c$ for $c \in [w, b)$ and $s \in S$.

Definition 5.1 \mathcal{C} is the class of preference orderings in \mathcal{P} that satisfy axioms C1-3.

Lemma 5.2 For a preference ordering \preceq in \mathcal{C} , represented by θ according to (4.2), the condition $f_s c \sim c$ implies that $c \in \text{range } f_s$, and is characterized by

$$f_s c \sim c \Leftrightarrow \min_{Q \in \Delta_s^+} E_s^Q f + \hat{\theta}_s(c, Q) = c \quad (5.1)$$

with

$$\hat{\theta}_s(c, Q) := \min_{T \in \Delta_s^+} \frac{\theta(c, TQ_s)}{T(s)}. \quad (5.2)$$

For θ minimal, $\hat{\theta}_s(c, Q) = \min\{c - E_s^Q f_s \mid f_s \in \mathcal{A}_s^c\}$ with $\mathcal{A}_s^c := \{f_s \in \mathcal{A}_s \mid f_s c \succeq c\}$.

The specification of Theorem 1 for \mathcal{C} can now be formulated as follows.

Theorem 2 A preference ordering \preceq in \mathcal{C} has a unique fixed point update (3.3) in state s , if it has a θ -representation so that

$$\text{the mapping } c \mapsto \hat{\theta}_s(c, Q) - c \text{ is strictly decreasing, for all } Q \in \Delta, \quad (5.3)$$

with $\hat{\theta}_s$ defined by (5.2). Then this unique fixed point update is regular, and represented by $\hat{\theta}_s$. Furthermore, then $\preceq_1 := (\preceq_s)_{s \in S}$ is sequentially consistent if

$$\min\{\theta(c, TQ_1) - E^T \theta_1(c, Q) \mid T \in \Delta\} = 0 \text{ for all } Q \in \Delta_1. \quad (5.4)$$

The conditions (5.3) and (5.4) are not only sufficient, but also necessary, for the minimal θ -representation (4.3) of \preceq .

So for minimal θ -representations of \preceq in \mathcal{C} , (5.3) is equivalent to axiom S3, and then axiom S2 amounts to (5.4).

6 Updating under Dynamic Risk Aversion

More explicit characterizations can be obtained under additional assumptions in terms of dynamic risk aversion. This is a somewhat stronger version of the notion of consistent risk aversion, introduced in RS16. Starting point is a given reference measure $P \in \Delta^+$ with respect to which risk aversion is defined. The idea is to impose, in addition to straightforward risk aversion with respect to P , also some dynamic properties in the same spirit. The vector $V_1(f)$ of certainty equivalents (ceqs) of f conditioned on S , is identified with an act in \mathcal{A} . As before, f stands for an arbitrary act in \mathcal{A} .

R1 (*Risk Aversion*) $f \preceq E^P f$.

R2 (*Consistent Risk Aversion, to S*) $f \preceq E^P V_1(f)$.

R3 (*Consistent Risk Aversion, from S*) $f \preceq E_1^P f$.

R4 (*Super-recursiveness*) $f \succeq V_1(f)$.

Definition 6.1 The class \mathcal{R} consists of preference orderings in \mathcal{S} that satisfy axioms R1-4.

Axiom R1 requires that risk premiums $V - E^P$ are nonnegative. It is easily verified that sequentially consistent updates then have the corresponding property,

$$f \preceq_1 E_1^P f \tag{6.1}$$

The axioms R2 and R3 reflect the intuition that ignoring risk premiums leads to higher values, over resp. the ‘stage’ towards S , and the stage from S . Axiom R3 is not used in the results below, but added for symmetry. Axiom R4 expresses that the overall risk premium should not exceed the aggregation of those towards and

from S , whereas the monotonicity axiom would require equality here.³ Notice that sequential consistency (axiom S1) directly follows from axiom R2 and R4, so axiom S2 and S3 are in fact redundant in the definition of \mathcal{R} .

Theorem 3 *Consider a preference ordering \preceq in \mathcal{C} that satisfies axiom R1, R2 and R4. It has a unique fixed point update in state s if and only if it has a θ -representation (4.2) with*

$$c \mapsto \theta(c, PQ_s) - cP(s) \text{ strictly decreasing, for all } Q \in \Delta.$$

If so, the fpu \preceq_s is regular, and represented by

$$\hat{\theta}_s(c, Q) := \frac{\theta(c, PQ_s)}{P(s)}.$$

The corresponding vector-update \preceq_1 is then sequentially consistent.

The proof in the appendix also contains a characterization of the θ -representations for each of the axioms R1–4.

6.1 The convex class and the reflection principle

Preference orderings \preceq induce a counterpart \preceq^* by the reflection principle

$$f \preceq^* g \Leftrightarrow -g \preceq -f. \tag{6.2}$$

This is a preference ordering on \mathcal{A}^* , the space of acts $\{-f \mid f \in \mathcal{A}\}$. Of particular interest is the case with symmetric outcome range $X = [-b, b]$, since then also \preceq^* is a preference ordering on \mathcal{A} . As \preceq^* shares many properties with \preceq , we call them twin preferences. Proofs of the list of claims below are left to the reader.

³In RS16, Chapter 7 it is described how this generalization amounts to the extension from recursive valuation per separate risk aversion level to a joint recursion in a range of levels.

All axioms A1-4, S1-3 and C2 hold for \preceq^* precisely when they hold for \preceq . The concavity axiom C1 itself reflects into convexity, and C3 yields $w_s c \prec^* c$ for constants $c \in (w, b]$. In a similar way, the inequalities in the anti-symmetric axioms R1-4 reverse for \preceq^* . The representation (4.2) transforms into

$$f \sim^* c \Leftrightarrow \max_{Q \in \Delta} E^Q f - \theta(-c, Q) = c.$$

The reflection principle commutes with fixed point updating, i.e., $(\preceq^*)_s = (\preceq_s)^*$, and we can simply write \preceq_s^* in the context of sequentially consistent updating.

Part II

Consistency without the monotonicity axiom

We now shift the focus from syntax to semantics. In Part I we have described an axiomatic framework with unambiguous updating that does not rely on the monotonicity axiom. This will serve as a platform to give our interpretation of preference reversals, and to develop concrete notions of choice consistency that are compatible with reversals. We do not strive for maximal generality. Our main purpose is to show that a normative interpretation is quite possible without the monotonicity axiom.

7 An elementary example

We start with an example to explain a subtle aspect of the interpretation of complete preference orderings. Consider the elementary lottery e with equal probability on

an outcome $+1$ and -1 , in dollars say, and an agent with preference ordering \preceq corresponding to expected utility with a strictly increasing concave utility function u . Then $e \sim -d$ for some constant $d > 0$. This is commonly phrased as “to the agent, e is indifferent to $-d$.” Consequently, the agent makes no difference between e and $-d$, it seems.

At the same time, however, $e \sim^* d$, with \sim^* denoting indifference under the twin preference \preceq^* defined by the reflection principle (6.2). Apparently this is a different form of indifference, but nevertheless a logical implication of \preceq (which, by the way, is in turn implied by \preceq^*). It is hence a matter of consistency that the agent has both indifferences in mind. Somehow e is indifferent to $-d$ as well to d . So, at closer inspection, the agent does make a difference between e and $-d$, since $-d \not\sim^* e$.

There is an obvious interpretation though, which is in fact the standard in bid-ask price models: $-d$ and $+d$ are resp. the agent’s bid and ask price for f . The underlying assumption is that obtaining an act f is identified with offering $-f$, which is plausible in a trading context. The agent puts two price tags on f , the ask price $V^*(f)$ and the bid price $V(f)$, both derived from one and the same preference ordering \preceq . Intuitively spoken, the preference ordering assigns a ‘thick’ value $[-d, d]$ to e , and ‘thin’ values c to constants c .

Consequently, the meaning of $e \sim -d$ should be more accurately phrased as “to the agent, *obtaining* e is indifferent to *obtaining* $-d$.” In other words, the value function V of \preceq is interpreted as a measure for the willingness to get, and V^* of \preceq^* for the willingness to keep, or the attractiveness to have.⁴

The crucial observation here is that the syntax of complete preference orderings does not impose that values are thin. The completeness of \preceq is respected, but not

⁴We could use the standard term ‘willingness to accept’ (WTA) here, but (i) accept is easily confused with obtaining, and (ii) in fact, it measures *unwillingness*, since higher outcomes means more attractive to keep. The so-called WTP-WTA-bias is discussed in Part III.

the completeness of its interpretation, as were it inducing ‘the’ value of e regardless the direction of trade. Such a value need not exist. One may try to argue, on normative grounds, that it should exist, or take the opposite standpoint, as we do. The mathematics leaves room for both. We will scrutinize both standpoints, first assuming a symmetric outcome range $[-b, b]$, to pay specific attention to the perhaps even more important case with nonnegative outcomes later on.

8 Absence of arbitrage

We first address absence of arbitrage opportunities, also called make-book or Dutch book opportunities. This is easy to obtain by a standard condition in bid-ask price modeling.

In this context, we take $V(f)$ as the agent’s price he is willing to pay for f . Axiom A2 already rules out the most direct form of arbitrage: the agent paying a positive amount for an act with only non-positive outcomes. Depending on the context, also series of acts f_1, \dots, f_K have to be excluded that have the same net effect, i.e., with non-positive sum yet positive sum of values. For $K = 2$ this amounts to excluding *round trip arbitrage*, with $f_2 = -f_1$. The following lemma formulates a few standard results that we use. We assume a symmetric outcome range $[-b, b]$.

Lemma 8.1 *A preference ordering \preceq in \mathcal{P} is arbitrage free when it satisfies axiom R1. It is free from round trip arbitrage iff $V \leq V^*$, i.e., for all $f \in \mathcal{A}$, if $f \sim c$, then $-f \preceq -c$. For concave preferences in \mathcal{C} this is the case if and only if for all $c \in [-b, b]$, there exists $P^c \in \Delta^+$ such that $\theta(c, P^c) = 0 = \theta(-c, P^c)$.*

It is essential for our interpretation that $V \leq V^*$, since it would be absurd to be willing to pay an amount c for f , and at the same time to offer it for less. We will adopt axiom R1 to guarantee this, and to keep at distance from arbitrage. This allows us to refer to V and V^* as resp. the upper and lower value induced by \preceq ,

following the terminology in e.g. Walley (1991). Updates inherit this ordering, i.e., $V_1 \leq V_1^*$, when they are sequentially consistent, cf. (6.1).

9 The monotonicity axiom revisited

Let us continue the quest for a normative interpretation of preferences \preceq in \mathcal{S} with sequentially consistent update \preceq_1 . We assume a symmetric outcome range $[-b, b]$, and have $V \leq V^*$ as consequence of axiom R1.

It is time for a direct confrontation with the monotonicity axiom (3.1). This is exactly the Sure Thing Principle, restricted to events $s \in S$,

$$f_s h \preceq g_s h \quad \Leftrightarrow \quad f_s h' \preceq g_s h'$$

provided that its compelling suggestion for updating is obeyed. It is also equivalent to backward recursiveness,

$$f_s \sim_s g_s \text{ on } S \quad \Rightarrow \quad f \sim g. \quad (9.1)$$

So the monotonicity axiom goes together with a reputable principle and a practical rule.

To proceed with the example on the elementary lottery e , we now direct the attention to an act f that contains e as a sub-lottery in a state $s \in S$, i.e., $f = e_s f$. We assume the same preference in state s as before, which is now assumed to be the update \preceq_s of some initial preference ordering \preceq in \mathcal{C} . So the agent considers the sub-lottery e as something with lower value $-d$ and upper value d .

The monotonicity axiom imposes that $f \sim (-d)_s f =: g$ (and also that $f \sim^* d_s f$). So the ceq of e should also be its replacement value. This means that the willingness to obtain f should depend on no other aspect of its sub-lottery e in state s than the willingness to obtain e again in that state; the same if it comes to offering f . Upper and lower values should be separately recursive. But why?

If the agent thinks through what it would mean to arrive in state s after acquiring f , it is with f_s already in his hands. The attractiveness to have f_s , $V_s^*(f_s)$, is part of the considerations to obtain f , naturally. Today's bid price may be influenced by tomorrow's ask price. In principle, all features of a sub-act may count for \preceq . Consequently, the replacement value of e in f , i.e., the constant r for which $e_s f \sim r_s f$, need not coincide with the ceq of e . Note that sequential consistency (axiom S1) only requires that e has replacement value $-d$ in $e_s(-d)$, but not in $e_s f$, so it gives room for non-consequentialist replacement values.

At closer inspection, monotonicity is not only debatable, it can be problematic. This is best explained in terms of market ask prices $V^*(f)$ with f resembling a financial asset, in line with RS13, Ex. 3.9. Monotonicity would imply that $V^*(f) = V^*(V_1^*(f))$, which means that a contract that pays out the ask price of f is equally expensive as f itself. Consequently, traders can avoid round trip costs by trading in such contracts rather than the asset itself. In the terminology of Gilboa (2015), this means that in such a context the monotonicity axiom not only loses its objectively rational (compelling) status, but that even its subjective rationality (being defensible) is questionable.

We conclude that in our interpretation the monotonicity axiom has lost the immediate normative appeal it has, when thickness of value is ignored. We accept that $f \succ g$, thus denying the axiom, and hence recursive valuation. Let us face the consequences.

10 Preference reversals and choice consistency

Relaxing the monotonicity axiom means accepting so-called dynamic inconsistency: a strict preference $f \succ g$ turning into $f_s \preceq_s g_s$ for all $s \in S$. This is exactly the case in the example above, when we allow that $f \succ g$. Since we defend this phenomenon,

we prefer to call it a (dynamic) preference reversal.⁵ The anomaly is perhaps felt most sharply in the setting of a dynamic choice problem, with f modified to f' by subtracting a small constant ε from e , so that still $f' \succ g$, yet $f' \prec_s g$. If the agent is offered the choice between f' and g , with the option to switch after s has obtained, he would choose f' first, *knowing* that he will switch to g in state s . The agent does not stick to his plan, it seems.

We want to defend the reversal, but not by giving up the stick-to-your plan principle. This is impossible, if that principle would be properly reflected by the monotonicity axiom. A closer inspection reveals that this is indeed not the case, from our viewpoint.

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. The most direct way to reflect this idea, tuned to a purely monetary setting in which the switch from f to g in s has net effect $g_s - f_s$, is to compare this net effect with 0. We therefore propose the following criterion for *Forward Monetary Choice Consistency*:

$$f \succ g \quad \Rightarrow \quad 0 \succ_s g_s - f_s \text{ for some } s \in S \quad (f, g, g - f \in \mathcal{A}) \quad (10.1)$$

A similar line of reasoning motivates the following static version.

MCC (*Monetary Choice Consistency*) $f \succ g \Rightarrow 0 \succ g - f \quad (f, g, g - f \in \mathcal{A})$.

The axiom 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

⁵The term preference reversal is also used for the combination $f \succ g$ yet $f \prec^* g$, as in e.g. Karni and Safra (1987). Such a ‘static’ reversal occurs whenever $\preceq \neq \preceq^*$, and this inequality has already been motivated.

g. It turns out that MCC implies (10.1), and also its backward analogue, *Backward Monetary Choice Consistency*:

$$f \succ_1 g \quad \Rightarrow \quad 0 \succ g - f \quad (f, g, g - f \in \mathcal{A}). \quad (10.2)$$

Theorem 4 *If \preceq in \mathcal{P} is monetarily choice consistent (axiom MCC), and has sequentially consistent update \preceq_1 , then also the forward and backward versions (10.1) and (10.2) are satisfied.*

For concave preferences, axiom MCC can be characterized as follows. We call $\theta(c, Q)$ strictly binding if there is an f in the interior of \mathcal{A} such that $E^Q f + \theta(c, Q) = c$. The lemma does not rely on symmetry of outcome ranges.

Lemma 10.1 *A concave preference in \mathcal{C} with θ -representation (4.2) satisfies axiom MCC if it has a θ -representation with $\theta(0, Q) = 0$ for all Q with $\theta(c, Q)$ strictly binding for some $c \in X$, and only if its minimal representation has this property.*

We remark that the MCC criterion has quite strong implications for $\theta(0, Q)$: it is either zero or ineffective. A milder criterion is obtained by restricting the MCC condition to pairs f, g with either f or g a constant. It can be shown that this is equivalent to the condition that for all Q , $\theta(c, Q)$ achieves its minimum in $c = 0$. We will not elaborate on this.

Our considerations so far lead to the following notion of consistency.

Definition 10.2 The class \mathcal{M} of *monetary consistent preferences* is the set of preferences in \mathcal{S} that satisfy axioms R1 and MCC.

The axioms S1-3 guarantee that updates in \mathcal{M} are well defined, consequentialist, and have the property that values $V(f)$ are in the range of the vector-update $V_1(f)$. Axiom R1 rules out arbitrage opportunities, and makes sure that value has nonnegative thickness $V^* - V$. Furthermore, preference reversals do not lead to

choice inconsistencies, when axiom MCC is imposed. We believe that this suffices to claim that a normative interpretation without the monotonicity axiom is quite possible, and even necessary when the latter leads to anomalies.

11 Twin consistent preferences

The criterion for choice consistency, axiom MCC, relies on a truly monetary setting in which the expression $g - f$ makes sense. This may hinder extensions to other settings, for instance with nonnegative outcomes, or in units of utility (utils). The following slightly different notions of choice consistency avoid the expression $g - f$.

FCC (*Forward Choice Consistency*) $f \succ g \Rightarrow V_1^*(f) \not\leq V_1(g)$.

BCC (*Backward Choice Consistency*) $f \succeq_1 g \Rightarrow V^*(f) \geq V(g)$.

So, where in the formalization of stick-to-your plan the monotonicity axiom compares $V_1(f)$ and $V_1(g)$, and (10.1) compares $V_1(g - f)$ and 0, the idea of FCC is to compare $V_1^*(f)$ and $V_1(g)$. A similar comparison applies to (10.2) and BCC (in which we switched to non-strictness for technical convenience).

This pair of notions can be characterized by a straightforward relaxation of recursiveness (9.1), that involves both \preceq and its twin \preceq^* .

TCC (*Twin Consistency*) $V(V_1(\cdot)) \leq V(\cdot) \leq V^*(V_1(\cdot)) \leq V^*(\cdot) \leq V^*(V^*(\cdot))$.

Theorem 5 *A preference ordering \preceq in \mathcal{S} satisfies axioms FCC and BCC if and only if it satisfies axiom TCC.*

Definition 11.1 The class \mathcal{T} of *twin consistent preferences* is the set of preferences in \mathcal{S} that satisfy axiom TCC.

It is easily verified that axioms S1 and S2 are redundant here, that axiom R4 holds true in \mathcal{T} , and that \mathcal{T} includes \mathcal{R} (even when axiom R3 is relaxed).

Our normative interpretation, in particular for \mathcal{T} , extends to setups which take starting point in nonnegative outcomes only. The point we distilled from the elementary example in Section 7 is less straightforward in this case, but perhaps therefore even more important. The twin preference \preceq^* then only applies to non-positive acts, and this domain has no overlap with that of \preceq , other than the constant 0. So \preceq^* on \mathcal{A} no longer derives from \preceq on \mathcal{A} , but depends on the way \preceq is extended to acts with symmetric outcome range $[-b, b]$. Then it is mathematically quite straightforward to insist on thin values for all acts, by taking $V(-f) = -V(f)$, so that $V^*(f) = V(f)$ on \mathcal{A} . In view of the elementary example, however, there is no reason to adopt this as a normative principle *per se*. In the contrary, our analysis suggests that a specification of the extension, be it in the form \preceq on $-\mathcal{A}$, or \preceq^* on \mathcal{A} , is an indispensable element in its normative interpretation. Value has thickness, $V^* - V$, whether one is aware of it or not, and this justifies to decouple consequentialist update values from non-consequentialist replacement values. We view this a fundamental aspect of complete preference orderings, also in non-monetary settings, as indicated in Part III.

Summarizing, we claim that regular preferences \preceq in \mathcal{P} , combined with a specification of \preceq^* in as far it does not derive from \preceq by reflection, admit a normative interpretation when axioms S3 and TCC are satisfied, despite the reversals they induce.

12 Application to the Allais paradoxes

The Allais paradox (Allais, 1953) involves four lotteries a_1, a_2, a_3, a_4 , with the following probability distributions on the outcomes 0,1 and 5 (\$mln):

$$\begin{aligned} a_1 \text{ has } p(1) = 1 & & a_2 \text{ has } p(0) = 0.01, p(5) = 0.10, p(1) = 0.89, \\ a_3 \text{ has } p(1) = 0.11, p(0) = 0.89 & & a_4 \text{ has } p(0) = 0.90, p(5) = 0.10. \end{aligned}$$

They can be represented in our setting as acts on an outcome space $\{\omega_1, \omega_2, \omega_3\}$, corresponding to the three outcomes of a_2 , in the obvious way. It has been well documented that many people prefer a_1 over a_2 , and a_3 over a_4 , contrary to the Independence Axiom, and hence to the Sure Thing Principle. The paradox reveals that the replacement value of the sublottery in $s := \{\omega_1, \omega_2\}$, is commonly perceived smaller than 1 in a_2 , and larger than 1 in a_3 . This sub-lottery, to which we refer as a_s , has $p(5) = 10/11$ and $p(0) = 1/11$. The monotonicity axiom is hence violated, for any choice of a consequentialist update in s .

The paradox resolves when replacement values are not identified with ceqs. The lotteries a_2 and a_3 contain the sub-lottery a_s that has some lower and upper value, and there is no reason to assume that replacing a_s by its lower value has no influence on the willingness to obtain.

A particularly simple idea is to consider worst expected values, with the given probabilities up- or downscaled by a at most a factor ten. It assigns a thick value $[0.05, 4.99]$ to a_s , replacement value 0.934 to a_s as part of a_2 , and 4.545 to a_s as part of a_3 . All these values follow from $V(f) = \min\{E^Q f \mid q_1 \geq 0.01, 0.001 \geq q_2 \geq 0.1, q_3 \geq 0.089\}$. Intuitively speaking, in a_2 conservatism concentrates on risk within the sub-lottery, and in a_3 on the risk of missing it, since that is most effective.

This preference ordering is in \mathcal{M} and \mathcal{T} (and also in \mathcal{C} and \mathcal{R}). The preference reversal for a_3, a_4 hence does not lead to a choice inconsistency, since axiom MCC and TCC hold true. Dutch book strategies have no chance, by Lemma 8.1.

So in our interpretation, the Allais paradox is not revealing an irrational preference reversal, but a sensitivity of replacement values of sub-lotteries for their context that perfectly makes sense.

13 Application to the Ellsberg paradox

The Ellsberg paradox concerns the following four gambles on the color of ball drawn from an urn with 30 red balls and 60 balls that are either yellow or black:

gamble A: \$100 when red gamble B: \$100 when black
 gamble C: \$100 when red or yellow gamble D: \$100 when black or yellow

Experiments indicate that most people strictly preferred A to B and D to C, which can be explained by the fact that A and D are bets with given probability on success while B and C are not. These preferences are not compatible with the Sure Thing Principle (and hence with expected utility), since $A \succ B$ then would imply that drawing black is considered as less likely than drawing red, while $C \prec D$ implies the opposite.

Ellsberg argues in his seminal paper (Ellsberg, 1961) that there are no compelling reasons to deem this irrational. He explains, among many other things, how the ordering can be explained in terms of preferences of the form

$$V(f) = \rho E^{Q_0} f + (1 - \rho) E^{Q_{\min}} f,$$

where Q_0 denotes the best estimate of the proportions of the colors, say $(1/3, 1/3, 1/3)$ by symmetry, ρ the degree of confidence in that best estimate, and Q_{\min} the worst case probability measure compatible with the experiment, i.e., with minimal probability on the winning state in f . He emphasizes that taking $\rho < 1$ is a matter of conservatism, not pessimism. While pessimism apparently falls short of rationalizing the preferences, conservatism does provide a consistent interpretation that can stand closer inspection.

We do not have much to add to this, other than confirming that these preferences are indeed choice consistent in the sense of axiom MCC as well as TCC. The classic choice inconsistency is just a reversal in lower values. This confirms Ellsberg’s view that the key to understand the paradox lies in recognizing that both inner and outer (or lower and upper) probabilities are in play. Ambiguity is not essential in this, in our interpretation, but probably enlarges the thickness of value that already exists under risk.

In summary, our framework supports the normative interpretation of the delicate aspects of decision making revealed by the Allais and Ellsberg experiments.

Part III

Related literature and conclusions

We first explain the background of this paper. Then we discuss related literature on two central topics: updating, and the discrepancy between buying and selling prices that we based on the reflection principle. We conclude by discussing some additional topics and closing remarks.

14 Background

The central notion in our framework, sequential consistency, has been introduced in Roorda and Schumacher (2007) (henceforth RS07) and further developed in mathematically more advanced settings in RS13; RS16.⁶ These papers describe axiomatic

⁶A common assumption in these papers is that value functions V are *translation invariant*, i.e., satisfy $V(f + c) = V(f) + c$. Sequential consistency then amounts to the criterion in axiom S1 restricted to $c = 0$.

frameworks with sequential consistency as basis for unique updating, applicable to both a pricing context, with $V(f)$ the bid price of f , as well as a regulatory context, with $E^P(f) - V(f)$ a required capital buffer against extreme losses.⁷ The fixed point update rule (3.3) closely relates to the conditionally consistent updating rule in RS07 and the refinement update introduced in RS13. It turned out that the fixed point update rule already was present in the literature on preference orderings, as discussed below.

For a general introduction to sequential consistency, and other forms of non-recursive, so-called *weak time consistency concepts* we refer to these papers and the references therein. That recursiveness is problematic in a regulatory context has been signaled in RS07, Ex. 8.8. Similar, yet less pronounced, concerns about recursive pricing are indicated in RS13, Ex. 3.9. The observed preference reversals described in these examples gave rise to further investigation at the level of complete preference orderings, of which the current paper is a reflection. As compared to the aforementioned papers, the new elements in this paper are the relaxation of translation invariance, the more extensive discussion of a normative interpretation, the results on choice consistency, and the connection with the Allais and Ellsberg paradox. Also, the updating formulas for quasi-concave value functions extend those in RS16 for ordinary concave valuations. The insights in Cerreia-Vioglio et al. (2011b) and Frittelli and Maggis (2011) provided the motivation and the tools for this extension.

⁷There is a continuous spectrum from pricing with small risk premiums (close to expected values) to so-called risk measures (much closer to worst-case), which explains that axiomatic frameworks in both domains are strikingly similar. For instance, the seminal paper Artzner et al. (1999) advocates coherent risk measures, which corresponds to MEU with trivial utility $u(x) = x$. Convex risk measures correspond to (ordinary) concave, translation invariant, monotone V , whereas convexity is not imposed for monetary risk measures, see Föllmer and Schied (2011).

15 On the monotonicity axiom

The monotonicity axiom (3.1) plays a prominent role in the literature on preference orderings. Three levels of definition can be discerned: in deterministic states, in unambiguous states, and in ambiguous states. We take the first level for granted, as reflected in axiom A2, but deny monotonicity at the other levels, which we treat at the same footing.

The standard monotonicity axiom in Anscombe-Aumann (AA)-settings has unambiguous states, with sub-lotteries identified with their probability distribution.⁸ This is often accepted as a basic rationality axiom, see e.g. Gilboa (2015). In our approach, the fpu rule (3.3) assigns to each roulette lottery its certainty equivalent, $V_s(f_s)$, in a consequentialist way, but its replacement value in the horse lottery (r such that $V(f) = V(r_s f)$) may depend on the roulette lotteries in other states. We only require, by axiom S2, that if all probability distributions have the same ceq c (if it comes to obtaining), then c is the ceq of the horse lottery as a whole (if it comes to obtaining).

Monotonicity with respect to ambiguous states often appears as consequence of the STP in Savage's setting, or as dynamic consistency requirement, when updating is considered with respect to a partitioning $\Pi = \{E_1, \dots, E_k\}$ of end states of horse lotteries in an AA-setting.

Normative objections against the monotonicity axiom in the literature concentrate on versions with ambiguous states (see Trautmann and Wakker (2015) and references therein). The objections confirm our viewpoint that replacement values are non-consequentialist by nature, but our line of reasoning is different, and does not rely on the distinction between risk and ambiguity. To our knowledge, the ax-

⁸The AA-setting can be approximated in our setting by taking Ω sufficiently large. State independence can be obtained by taking $\Omega = S \times \Omega'$. Monotonicity in final outcomes (axiom A2) amounts to first-order stochastic dominance under law-invariance.

omatic alternative we propose, in particular the revision of the definition of choice consistency, brings a new element to these considerations, as further discussed in the next section.

A comparison between the Sure Thing Principle (STP) and Axiom S2 makes clear that the class of regular preferences for which the fpu indeed produces a sequentially consistent update goes far beyond what the monotonicity axiom tolerates. Axiom S2 is weaker than the Comonotonic STP (Wakker, 2010), since it only imposes a condition on acts with the same value in each state S . We refer to RS16 for concrete descriptions of the extra freedom, allowing for jointly recursive evaluation at a range of risk aversion levels. How this relates to the empirical findings on violations of the monotonicity axiom is a topic of future research.

16 On the fixed point update

The fixed point update rule (3.3) is not new. It 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 (3.3). In Siniscalchi et al. (2001) it appears as constant-act dynamic consistency, and is interpreted as fixed point criterion. Its close connection with the Generalized Bayesian Rule in Walley (1991) and the Full Bayesian Updating Rule in Jaffray (1994) is well understood for the Gilboa-Schmeidler framework with Multiple Priors, also called Maxmin Expected Utility (MEU), see Pires (2002). In Eichberger et al. (2007) this connection is addressed at the level of capacities, and applied to the class of Choquet Expected Utility (CEU), also where it lies outside MEU (see also Horie (2007) for a correction).

We see the following points of contribution with respect to the existing results on update rules. We provided a further underpinning of the fpu rule, as the only candidate for a sequentially consistent update (axiom S1), which led to an additional

requirement (axiom S2) to ensure that the rule indeed produces a consistent update. Furthermore, we extended its scope considerably. The rule is well defined for any preference ordering that is regular (axioms A1-4, class \mathcal{P}) and is c -Sensitive (axiom S3). So, under only a few technical conditions, the fpu produces an update for any complete preference ordering, and sequential consistency rules out any alternative. In particular, it defines updates for the quite general class of Uncertainty Averse Preferences (Cerreia-Vioglio et al., 2011a), which covers all known models that are extensions of MEU, and also Smooth Ambiguity Preferences (Klibanoff et al 2005). For instance, Variational Preferences (Maccheroni et al 2006) are those with θ -representations (4.2) independent of c . The specification of Thm. 2 to the case with $\theta(c, Q) \in \{0, \infty\}$ and independent of c , provides the criterion for MEU, and hence for convex capacities in CEU.⁹ Moreover, utility functions with convex and concave segments are included, so that the fpu also applies to Cumulative Prospect Theory (CPT) (Tversky and Kahneman, 1992). For all these models, fpu is hence the adequate update principle if sequential consistency is adopted.

Doubts have been expressed that such a universal update principle can exist outside expected utility models (Wakker, 2010; Machina and Viscusi, 2013). Under the classic definition of choice consistency, which amounts to the exclusion of preference reversals, indeed no update rule seems satisfactory. If $f \succ g$ yet $f \prec_1 g$ has to be avoided, there are basically two remedies. The forward induced choice ($f \succ g$) is given priority, and the update rule is adjusted to avoid reversals, as proposed in Hanany and Klibanoff (2007), but this is not consequentialist. Alternatively, backward induced conditional choices ($f \prec_s g$) are ‘folded back’ to initial choices,

⁹This criterion has been interpreted as a junctedness condition in RS07, which imposes that each set of conditional probabilities $P^i(\cdot|E_i)_{i \in \mathbb{N}}$ with $\theta(P^i) = 0$ that occurs, also occurs as the set of conditional probabilities of one ‘junct’ P^* with $\theta(P^*) = 0$ (see also RS16, Section 4 for the concave case). It shows the difference with the rectangularity condition implied by the STP for Π , as in Epstein and Schneider (2003).

according to the principle of *consistent planning* (Siniscalchi, 2009) and *behavioral consistency* (Karni and Safra, 1990), which also has problematic aspects (Machina, 1989). Our interpretation made us find a third way out, that of rethinking the very definition of choice consistency. We achieve the stick-to-your-plan principle under uncomplicated updating, by recognizing that plans are sticky, so to speak.

17 On the reflection principle and bid-ask prices

One of the cornerstones of our framework is the observation that acts have not just one price, but (at least) two. This is well recognized in the literature, in several ways, under the heading of e.g. first order risk aversion, endowment effect, the WTP-WTA bias, and the law of two prices. Many aspects of our framework are already treated in these branches of literature, at a deeper level, and in a much more advanced mathematical and economic setting. An overview is far beyond our scope, and we only discuss a few representative examples, with a focus on the given interpretation of bid-ask spreads.

The reflection principle is a standard way to relate bid and ask prices, in particular in monetary settings. It is used for instance in *conic finance*, introduced in Madan and Cherny (2010) as a new way to model markets with bid-ask spreads.¹⁰ Also the way in which loss aversion and probability weighting induce bid-ask spreads, as described in (Wakker, 2010, Ex. 6.6.1 and 9.3.2),¹¹ is in line with the reflection principle. For instance, in RDU, the reflection principle amounts to reversed rank order for ask prices. In fact, the findings in Birnbaum and Stegner (1979) already

¹⁰The term conic refers to cones as the acceptance set of so-called coherent risk measures, which have been introduced by Artzner et al. (1999).

¹¹Loss aversion in CPT is a form of first-order risk aversion, which corresponds to utility functions with a kink in a reference point, commonly at zero, cf. (Segal and Spivak, 1990).

point in this direction. It links bid-ask spreads to so-called *configural weighting*, in the context of estimating used car prices. They find that “Judges instructed to take the buyer’s point of view gave greater weight to the lower estimate, whereas judges who identified with the seller placed a greater weight on the high estimate,” and they emphasize the point that it is the same cautiousness that results in different prices for opposite directions of trade.

To discern two prices for the same thing is hence by no means new, nor to base it on reflection, but we think that the consequences for interpretation and rationalization have not yet been fully recognized. For instance, in conic finance, still recursiveness is commonly imposed in bid and ask prices separately (Madan, 2016), so that the induced (dynamic) preference reversals are avoided. However, as we have indicated in Section 9, this may lead to a market in which round trip costs can always be avoided. An exception is Bielecki et al. (2013), which applies a notion of weak time consistency (their definition D7) to conic finance that corresponds to sequential consistency; the special role of the fpu is however not addressed, nor the idea of a joint recursion involving bid and ask prices. An interesting point of future research is to test for such a joint recursion on the basis of market prices.

Empirical research has led to extensive descriptions and interpretations of the discrepancy between selling and buying prices, or closely related notions. Common terms in this literature, like *WTP-WTA bias* (Machina and Viscusi, 2013, Chapter 4), *endowment effect* (Kahneman et al., 1990), (static) *preference reversal* (Karni and Safra, 1987), *failure of procedure invariance* (Tversky and Thaler, 1990) indicate that two values for the same thing is primarily viewed as irrational by nature. Our results give reason to rethink the normative content of these findings, and provides a general framework for further developing normative bid-ask models. Issues observed in eliciting certainty equivalents, as described in (Machina and Viscusi, 2013, Chapter 4), may be clarified by a careful distinction between replacement

values and certainty equivalents, both for obtaining and offering. The notions of sequential consistency and choice consistency we proposed suggest new empirical questions on dynamic choice.

Intuitively speaking, our analysis emphasizes the importance of keeping ‘looking through’ sub-acts, rather than perceiving them as represented by just one ‘thin’ value and then concentrate on the weight of that value. We do agree that such a weight is configural, but it need not be a function of the ceqs of sub-acts alone.

18 Other topics

We have proposed two notions of choice consistency in Section 10, namely MCC and TCC. To our knowledge, the interpretation of MCC as static requirement to ensure the dynamic versions ((10.1) and (10.2)) is new. The criterion for MCC in Lemma 10.1 exhibits similarities with first-order risk aversion, that deserve further investigation. We did not find a reference that relates the elementary conditions in the second notion, TCC, to choice consistency.

There is an interesting link with how Bewley’s incomplete preferences are used in Gilboa (2015) to model objective rationality. A given preference \preceq induces an incomplete preference ordering \preceq^i , with $f \preceq^i g$ if and only if (even) the ask price of f is at or below the bid-price of g , i.e., $V^*(f) \leq V(g)$. Forward choice consistency (axiom FCC) requires that when $f \preceq_1^i g$, it cannot be that $f \succ g$. If we interpret \preceq_1^i as the ‘objectively rational’ part of \preceq_1 , axiom FCC can be phrased as “a preference ordering should respect the objectively rational consequences of its update”.

We relied on the reflection principle (6.2) to justify the disconnection of ceqs and replacement values. Although we believe that it is rather compelling in a monetary setting, it may require modification in other contexts, especially when the negative of outcomes is not well defined. For instance, one can first express outcomes in

utils, and then identify offering x utils with obtaining $-x$ utils. When the utility function u is continuous and strictly monotone, the axioms for \mathcal{S} each amount to the same for \preceq and the transformed ordering $\tilde{\preceq}$ on acts $u \circ f$. This indicates that essence of our interpretation is also relevant in non-monetary settings.

19 Concluding remarks

We have proposed a normative interpretation for preferences with reversals, in which consequentialist updating is combined with choice consistency. Our key observation is that there is a twin side inherent in preference orderings that cannot be locked up by the Sure Thing Principle, but requires the extra space offered by the Equal Level Principle (axiom S2). The proposed notions of choice consistency are tuned to this viewpoint, and pull out the sting of the Allais and Ellsberg paradoxes. Updating remains remarkably simple in our framework.

This extension brings an extra dimension to the many concepts and techniques rooted in the monotonicity axiom for preference orderings, such as risk neutral valuation, dynamic programming, sub-game perfectness, and exponential discounting. It invites less monotone ways to break time and value into pieces, accepting preference reversals without a sigh.

20 Appendix

20.1 Proof of Theorem 1

We first prove that (3.3) defines a unique update \preceq_s under S3, for each $s \in S$. Let V denote the (normalized) value function of \preceq . Consider, for given $f_s \in \mathcal{A}_s$, the mapping $\rho : c \mapsto V(f_s c)$ on the domain $\text{range}(f_s) =: [l, r]$. Since V is continuous and monotone, 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 (3.3). Axiom S3 guarantees that such c' is unique, and hence that \preceq_s is uniquely determined by (3.3). This means that \preceq_1 is indeed unambiguously defined by (3.3).

This proves the if-part of the first claim of the theorem. The only if-part is obvious from the formulation of S3.

Regularity of \preceq_s , under axiom S3, follows straightforwardly 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 (3.3) for f_s^k , any converging subseries $(c_k)_{k \in \mathcal{I} \subset \mathbb{N}} \rightarrow c'$ yields $V(f_s c') = c' \in \text{range}(f_s)$, by continuity of V ; so c' must be the unique solution of the rhs in (3.3), and hence the full series $(c_k)_{k \in \mathbb{N}}$ is converging to c' .

To see that \preceq_1 defined by (3.3) is sequentially consistent if \preceq satisfies axiom S2, consider $f \in \mathcal{A}$ with $f \sim_1 c$. Then (3.3) implies that for all $s \in S$, $f_s c \sim c$ with $c \in \text{range}(f_s)$, and by axiom S2, $f \sim c$, so that axiom S1 follows.

It remains to show, under axiom S3, that if \preceq has a regular sequentially consistent update \preceq_1 , then \preceq must satisfy axiom S2. 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 $f_s c' \sim c'$ by axiom S1, while also $f_s c \sim c$ by assumption, and axiom S3 implies that $c' = c$. Since $s \in S$ was arbitrary, $f_s \sim_s c$ for all $s \in S$, and, again by axiom S1, indeed $f \sim c$.

20.2 Proof of Lemma. 4.1

We prove the lemma in terms of R -representations (4.4), in line with Cerreia-Vioglio et al. (2011b), using the 1-to-1 correspondence (4.5), which is easily verified. Note that $m \mapsto R(m, Q)$ is the inverse of the strictly increasing mapping $c \mapsto c - \theta(c, Q)$. The properties (i)-(iii) translate for R into (i') R is strictly increasing in m , (ii') $\min_Q R(m, Q) = m$, and (iii') R is equi-continuous in m , i.e., for all $\varepsilon > 0$, there

exists $\delta > 0$ such that for all $Q \in \Delta$, $|m' - m| < \delta$ implies $|R(m', Q) - R(m, Q)| < \varepsilon$.

For the if-claim, we have to prove that V in (4.4) is monotone, normalized, continuous, and quasi-concave if R satisfies (i')-(iii'). Monotonicity is obvious, and (ii') implies that V is normalized, i.e., that $V(c) = c$. That V is quasi-concave follows from (i') and the inequality $E^Q(\lambda f + (1 - \lambda)g) \geq \min\{E^Q f, E^Q g\}$. It remains to show that V is continuous.

Consider a series $f^n \rightarrow \bar{f}$ and define $c_n := V(f^n)$ and $c := V(\cdot)$. We have to prove that $c_n \rightarrow c$. From (4.4) it follows that for all n , $R(E^Q f^n, Q) \geq c_n$ for all Q , and that equality is reached for some $Q^n \in \Delta$. From continuity of $R(m, Q)$ in m it follows that $R(E^Q \bar{f}, Q) \geq \limsup c_n =: \hat{c}$, for all $Q \in \Delta$, and hence $V(\bar{f}) \geq \hat{c}$. On the other hand, $\liminf R(E^{Q^n} \bar{f}, Q^n) \leq \liminf c_n =: \bar{c}$, and by equi-continuity of R hence $V(\bar{f}) \leq \bar{c}$. It follows that $V(\bar{f}) = \hat{c} = \bar{c} = \lim c_n = c$.

For the only-if claim under axiom C2, consider a concave preference ordering in \mathcal{P} . Let R be given by (4.6), which corresponds to θ in (4.3), and define $\bar{V}(f) := \min_{Q \in \Delta} R(E^Q f, Q)$. We have to prove that $\bar{V} = V$, and that R satisfies (i')-(iii'). It is obvious that R is minimal.

That $\bar{V} = V$ is a standard duality result: $\bar{V} \geq V$ by construction, for any \preceq , since for f with $V(f) = c$, $R(E^Q f, Q) \geq c$ on Δ , and hence $\bar{V}(f) \geq c$. By a standard separating hyperplane argument, which relies on the convexity of the upper contour sets \mathcal{A}^c (4.1), equality follows when \preceq is concave.

To derive (i'), notice that $R(m, Q)$ is non-decreasing in m , since V is monotone. If it were not strictly increasing, say $R(m, Q) = R(m + d, Q)$ for some m and Q , then there would exist an f with $V(f) = R(m, Q)$, by (4.6), and hence with $V(f + d) \leq R(m + d, Q) = V(f)$, which violates axiom C2.

Property (ii') follows from $V(c) = c$. This implies that $R(c, Q) \geq c$ on Δ , and also that $V(c) = R(E^Q c, Q) = R(c, Q) = c$ for some $Q \in \Delta$, by (4.4) for $f = c$. So $\min_{Q \in \Delta} R(c, Q) = c$, which is (ii').

Finally, (iii') follows from continuity of V . By continuity of V , for all $\varepsilon > 0$ there exists $\delta > 0$ such that $|V(f) - V(g)| < \varepsilon$ for all $|f - g| < \delta$, with $|\cdot|$ the sup-norm. Then also $R(m + d, Q) - R(m, Q) < \varepsilon$ for all $Q, m, d < \delta$, since

$$\begin{aligned} R(m + d, Q) &= \max\{V(f) \mid E^Q f \leq m + d\} \\ &= \max\{V(g + h) \mid E^Q g \leq m, |h| \leq d\} \\ &\leq \max\{V(g) \mid E^Q g \leq m\} + \varepsilon = R(m, Q) + \varepsilon \end{aligned}$$

Note that indeed ε is independent of m and Q .

20.3 Proof of Lemma 5.2

The preference ordering \preceq has a θ -representation (4.2), since it satisfies axiom C1 and C2 by assumption. Axiom C3, restricted to a given state $s \in S$, can be characterized in terms of (any) θ -representation of \preceq by the criterion

$$\theta(c, Q) = 0 \text{ with } c < b \Rightarrow Q(s) > 0. \quad (20.1)$$

This can be seen as follows. If the criterion does not hold true, for some c, Q, s , there exists a $Q, c < b$ with $Q(s) = 0$ and $\theta(c, Q) = 0$, and then $E^Q b_s c + \theta(c, Q) = c$ and hence $b_s c \sim c$. Conversely, if $b_s c \sim c$ for some $c < b$, there exists a $Q \in \Delta$ that achieves the minimum in (4.2) for $f = b_s c$, which is $E^Q b_s c + \theta(c, Q) = c$. This Q must have $Q(s) = 0$ and $\theta(c, Q) = 0$, and hence (20.1) does not hold true.

Hence, for (any) θ that represents \preceq , we have, with $\Delta_s^+ := \{Q \in \Delta \mid Q(s) > 0\}$,

$$\begin{aligned}
f_s c \sim c &\Leftrightarrow \min_{Q \in \Delta} E^Q f_s c + \theta(c, Q) = c \\
&\Leftrightarrow \min_{Q \in \Delta_s^+} E^Q f_s c + \theta(c, Q) = c \\
&\Leftrightarrow \min_{Q \in \Delta_s^+} E_s^Q f + \frac{\theta(c, Q)}{Q(s)} = c \\
&\Leftrightarrow \min_{T, Q \in \Delta_s^+} E_s^Q f + \frac{\theta(c, TQ_s)}{T(s)} = c \\
&\Leftrightarrow \min_{Q \in \Delta_s^+} E_s^Q f + \hat{\theta}_s(c, Q) = c
\end{aligned}$$

That $f_s c \sim c$ only admits solutions $c \in \text{range}(f_s)$, is obvious from the third line above, combined with the fact that $\min\{\theta(c, Q) \mid Q \in \Delta^+\} = 0$, by (20.1).

Note that $\hat{\theta}_s(c, RQ_s)$ is independent of $R \in \Delta_s^+$, so that we can write $\hat{\theta}_s(c, Q_s)$ with Q_s the conditional probability of $Q \in \Delta_s^+$ in s . We show that when θ is minimal, $Q_s \mapsto \hat{\theta}_s(c, Q_s)$ is lower semi-continuous (l.s.c.) and convex in Q_s , for all $c \in X$. By a standard duality result, then $\hat{\theta}_s(c, \cdot)$ must be the minimal representation of \mathcal{A}_s^c , which is closed and convex, so that the last claim follows.

To establish l.s.c., consider a series $Q_s^n \rightarrow \bar{Q}_s$, and take T^n such that $\hat{\theta}_s(c, Q_s^n) = \theta(c, T^n Q_s^n)/T^n(s)$. Axiom C3 implies that $T^n(s)$ is bounded away from zero, as follows. Note that $\hat{\theta}_s(c, \cdot) \leq c - w$, hence bounded, since in (5.2) we can take T with $T(s) = 1$, and $\theta(c, \cdot) \leq c - w$ for minimal θ , by (4.3). So if a sub-series has $T^n(s) \rightarrow 0$, then it also must have $\theta(c, T^n Q_s^n) \rightarrow 0$, and hence, because θ is l.s.c., have a limit point T' with $\theta(c, T') = 0$ and $T'(s) = 0$, precisely what is excluded by axiom C3, see (20.1). So $T^n(s)$ must be bounded away from zero. Consequently, there must exist $\bar{T} \in \Delta_s^+$ such that $\liminf \theta(c, T^n Q_s^n) = \theta(c, \bar{T}\bar{Q}_s)$, and hence

$$\liminf \hat{\theta}_s(c, Q_s^n) = \liminf \theta(c, T^n Q_s^n)/T^n(s) \geq \theta(c, \bar{T}\bar{Q}_s)/\bar{T}(s) \geq \hat{\theta}_s(c, \bar{Q}_s),$$

where the first inequality follows from l.s.c. of θ , and the second by definition of $\hat{\theta}_s$.

To derive convexity, consider $\tilde{Q} = \lambda Q + (1 - \lambda)Q'$ for some $\lambda \in (0, 1)$. We have to show that $\hat{\theta}_s(c, \tilde{Q}) \leq \lambda \hat{\theta}_s(c, Q) + (1 - \lambda) \hat{\theta}_s(c, Q')$. By definition of $\hat{\theta}_s$, the convex combination is equal to

$$\lambda \frac{\theta(c, TQ_s)}{T(s)} + (1 - \lambda) \frac{\theta(c, T'Q'_s)}{T'(s)}$$

for some $T, T' \in \Delta_s^+$. This can be rewritten as

$$\frac{\theta(c, \tilde{T}\tilde{Q}_s)}{\tilde{T}(s)} \text{ with } \tilde{T} = \mu T + (1 - \mu)T', \mu = \frac{\lambda T'(s)}{\lambda T'(s) + (1 - \lambda)T(s)}.$$

By definition of $\hat{\theta}_s$, it follows that this dominates $\hat{\theta}_s(c, \tilde{Q})$.

20.4 Proof of Theorem 2

To derive that the criterion (5.3) is sufficient for uniqueness of fixed point updates (and hence implies axiom S3), suppose $f_s c \sim c$ and $f_s d \sim d$ for some $c < d$. From the first indifference, it follows by (5.1) that there exists Q^* with $E_s^{Q^*} f + \hat{\theta}_s(c, Q^*) - c = 0$. But then, since also $f_s d \sim d$, $E_s^{Q^*} f + \hat{\theta}_s(d, Q^*) - d \geq 0$, and hence, contrary to (5.3),

$$\hat{\theta}_s(c, Q^*) - c \leq \hat{\theta}_s(d, Q^*) - d. \quad (20.2)$$

To prove the necessity of (5.3) when θ is minimal, suppose (20.2) holds true. By Lemma 5.2, there exists f_s with $f_s d \sim d = E_s^{Q^*} f_s + \hat{\theta}_s(d, Q^*)$, while, by (20.2), $f_s c \preceq c$. So axiom S3 is not satisfied, and hence fixed point updates are not unique.

That the unique fpu is regular, has already been proved in Theorem 1, and that it is represented by $\hat{\theta}_s$ is obvious.

The criterion (5.4) is essentially the same as in RS16, Corollary 4.3, and the proof that it characterizes sequential consistency, is entirely analogous, though technically far less complicated due to the simplicity of our mathematical setting.¹² Sufficiency

¹²On the other hand, in RS16 only *translation invariant* value functions are considered, with $\mathcal{A}^c = \mathcal{A}^0 + c$, cf. (4.1), so that only the case $c = 0$ has to be addressed. In our setting, for a fixed c we can define $\tilde{\mathcal{A}}^d := \mathcal{A}^c + d - c$, which is translation invariant, and apply the proof in RS16 to $\tilde{\mathcal{A}}^0$.

of (5.4) is straightforward, and its necessity is derived from a separating hyperplane argument.¹³

20.5 Proof of Theorem 3

As an additional result, we first characterize the axioms R1–4.

Lemma 20.1 *Let representations θ and θ_1 be given that represent resp. \preceq and \preceq_1 , by (4.2), and let V denote the value function of \preceq . The pair \preceq, \preceq_1 satisfies*

1. *axiom R1 if $\theta(c, P) = 0$*
2. *axiom R2 if $\theta(E^P c_1, PQ_1) \leq E^P \theta_1(c_1, Q)$*
3. *axiom R3 if $\theta(c, Q'P_1) \leq \theta(c, Q'Q_1)$*
4. *axiom R4 if $\theta(V(c_1), Q'Q_1) - V(c_1) \geq E^{Q'}(\theta_1(c_1, Q) - c_1)$*

for all $c \in X$, $c_1 \in X^{\#S}$, $Q' \in \Delta$, $Q \in \Delta^+$. Conversely, each of the four conditions is implied by resp. axiom R1–4 for minimal θ and θ_1 .

Proof of the auxiliary lemma

On axiom R1. sufficiency of the criterion is obvious from (4.2), necessity from (4.3).

On axiom R2. For the sufficiency of the criterion, consider f with $V_1(f) = c_1 = E_1^{\tilde{Q}} f + \theta_1(c_1, \tilde{Q})$, for some $\tilde{Q} \in \Delta^+$. Then

$$E^P V_1(f) = E^{P\tilde{Q}_1} f + E^P \theta_1(c_1, \tilde{Q}) \geq E^{P\tilde{Q}_1} f + \theta(E^P c_1, P\tilde{Q}_1),$$

and hence $f \preceq E^P c_1$. For its necessity, suppose it does not hold true for minimal θ , i.e., for some \tilde{c}_1, \tilde{Q} ,

$$\theta(E^P \tilde{c}_1, P\tilde{Q}_1) > E^P \theta_1(\tilde{c}_1, \tilde{Q}) \tag{20.3}$$

¹³Applied to two disjoint convex sets, denoted as \mathcal{Y} and \mathcal{Z} . In our setting, for given $c \in X$, we can take $\mathcal{Z} := \{(Q, t) | t \geq \theta(c, Q)\}$ and $\mathcal{Y} := \{(Q, t) | t \leq E^Q \theta_1(c, Q^*)\}$, where Q^* denotes a probability measure for which (5.4) is not satisfied.

Define $\tilde{c} := E^P \tilde{c}_1$. Since θ is minimal, there exists $\tilde{f} \in \mathcal{A}$ with

$$V(\tilde{f}) = E^P \tilde{c}_1 = E^{P\tilde{Q}_1} \tilde{f} + \theta(\tilde{c}, P\tilde{Q}_1) = \tilde{c} \quad (20.4)$$

(note that generally not $V_1(\tilde{f}) = \tilde{c}_1$). But then, contrary to axiom R2, $E^P V_1(\tilde{f}) < \tilde{c}$, which can be derived as follows. From (20.3) and (20.4), $E^P [E_1^{\tilde{Q}} \tilde{f} + \theta_1(\tilde{c}_1, \tilde{Q})] < \tilde{c}$, so there must exist $s \in S$ with $E_s^{\tilde{Q}} \tilde{f} + \theta_s(\tilde{c}_s, \tilde{Q}) < \tilde{c}_s$. Hence, by property (ii) in Lemma 4.1, there exist $\delta > 0$ so that $E_s^{\tilde{Q}} \tilde{f} + \theta_s(\tilde{c}_s - \delta, \tilde{Q}) = \tilde{c}_s - \delta$. Define $t_1 := \tilde{c}_1 - 1_s \delta$. Using an alternative expression for V_1 , namely

$$V_1(f) = \min\{E_1^Q f + \theta_1(c_1, Q) \wedge c_1 \mid c_1 \in X^{\#S}, Q \in \Delta\},$$

it follows that

$$E^P (V_1(\tilde{f})) \leq E^P (E_1^{\tilde{Q}} \tilde{f} + \theta_1(t_1, \tilde{Q}) \wedge t_1) < E^P \tilde{c}_1 = \tilde{c}.$$

On axiom R3. For the sufficiency of the criterion, consider f with $E_1^P f =: g \sim c$. By (4.2) there exists \tilde{Q}' with $E^{\tilde{Q}'} g + \theta(c, \tilde{Q}') = c$. The criterion implies that $\theta(c, \tilde{Q}') \geq \theta(c, \tilde{Q}' P_1)$, so $V(f) \leq E^{\tilde{Q}' P_1} f + \theta(c, \tilde{Q}' P_1) \leq c$.

For its necessity, suppose the criterion does not hold true for minimal θ , i.e, there exists c, \tilde{Q}' with $\theta(c, \tilde{Q}' P_1) > \theta(c, \tilde{Q}')$. Since θ is minimal, there exists $f \sim c$ with $E^{\tilde{Q}' P_1} f + \theta(c, \tilde{Q}' P_1) = c$. But then $E_1^P f \prec c$, as $E^{\tilde{Q}'} (E_1^P f) + \theta(c, \tilde{Q}') < c$.

On axiom R4. We first derive an expression for the minimal representation of $\check{V} := V(V_1(\cdot))$. Define $\check{\theta}$ by $\check{\theta}(c, Q) = \max\{c - E^Q f \mid \check{V}(f) \geq c\}$. Note that \check{V} is continuous and concave, so that (4.3) indeed applies. Rewrite, with θ_1 minimal,

$$\begin{aligned} \check{\theta}(c, Q) &= \max\{c - E^Q f \mid V(V_1(f)) \geq c\} \\ &= \max_{\{c_1 \mid V(c_1) = c\}} \max\{c - E^Q f \mid V_1(f) \geq c_1\} \\ &= \max_{\{c_1 \mid V(c_1) = c\}} \max\{c - E^Q c_1 + E^Q(c_1 - E_1^Q f) \mid V_1(f) \geq c_1\} \\ &= \max_{\{c_1 \mid V(c_1) = c\}} \{c - E^Q c_1 + E^Q \theta_1(c_1, Q)\} \end{aligned}$$

Axiom R4 amounts to the requirement that $\check{\theta} \leq \theta$, and this is exactly the criterion.

End of the proof of the auxiliary lemma

To prove the theorem, notice that the criterion for R4, with $c_1 = c \in X$, yields $\theta(c, Q'Q_1) \geq E^{Q'}\theta_1(c, Q)$. Combined with R1, this implies $\theta_1(c, P) = 0$, and with R2 it yields $\theta(c, PQ_1) = E^P\theta_1(c, Q)$. Taking Q of the form PQ_s , i.e., only differing from P in s , gives $\theta(c, PQ_s) = E^P\theta_1(c, PQ_s) = P(s)\theta_s(c, PQ_s)$. This implies that the minimum in (5.2) is achieved for $T = P$. All claims follow.

20.6 Proof of Lemma 8.1

For the first claim, consider f_1, \dots, f_K with $\sum_k f_k \leq 0$. By axiom R1, we have $V \leq E^P$, and hence $\sum_k V(f_k) \leq \sum_k E^P f_k = E^P \sum_k f_k \leq 0$, which we had to show.

By definition, a round trip arbitrage opportunity is the existence of $f \in \mathcal{A}$ with $V(f) + V(-f) > 0$. The second claim follows from $V(-f) = -V^*(f)$, cf. (6.2).

The condition in the last claim implies that the two convex spaces \mathcal{A}^c and $-\mathcal{A}^{-c}$ (see (4.1)) are separated by $\{f \mid E^{P^c} f = c\}$ (not strictly, because c is in their intersection). This implies $V \leq V^*$. Conversely, $V \leq V^*$ implies that there is a separating hyperplane between both aforementioned spaces. This takes the form $\{f \mid E^Q f = c\}$ for some $Q \in \Delta$. By axiom C3, $Q \in \Delta^+$, cf. (20.1), so we can take Q as P^c , and the last claim follows.

20.7 Proof of Theorem 4

The implication (10.1) follows from $f \succ g \Rightarrow 0 \succ g - f \Rightarrow 0 \not\preceq_1 g - f$, by resp. MCC and (3.2). The implication (10.2) follows similarly, from $f \succ_1 g \Rightarrow 0 \succ_1 g - f \Rightarrow 0 \succ g - f$, where for the first implication we now used the fact that a sequentially consistent update \preceq_1 inherits the static property MC (in each $s \in S$) from \preceq . To see this, assume $f_s \succ_s g_s$. Then, by axiom S3, $f_s c \succ g_s c$ for $c \sim f$, and hence $g_s 0 - f_s 0 \prec 0$, by MCC for \preceq , so $g_s - f_s \prec_s 0$, again by S3.

20.8 Proof of Lemma 10.1

Observe that MCC is not weakened if we impose it for f, g with $f, g, g - f$ in the interior of \mathcal{A} , since V is continuous. So MCC is the condition that $f + h \succeq f$ for all $f, f + h$ in the interior of \mathcal{A} with $h \succeq 0$.

To derive that the criterion is necessary for MCC, suppose there exists \tilde{Q} with (i) $\theta(0, \tilde{Q}) > 0$ and (ii) $\theta(c, \tilde{Q})$ strictly binding, i.e., with $E^{\tilde{Q}}\tilde{f} + \theta(c, \tilde{Q}) = c$ for some $\tilde{f} \sim c$ in the interior of \mathcal{A} . From (i), and minimality of θ , it follows that there exists $h \sim 0$ with $E^{\tilde{Q}}h + \theta(0, \tilde{Q}) = 0$, hence $E^{\tilde{Q}}h < 0$. By (ii), $\tilde{f} \sim c \succ \tilde{f} + \lambda h$ for all $\lambda \in (0, 1]$, and $\tilde{f} + \lambda h$ is in the interior of \mathcal{A} for sufficiently small λ , contrary to MCC.

To derive sufficiency, consider $c \sim g \prec f$. We can assume that g is in the interior of \mathcal{A} (otherwise add a sufficiently small constant to g). So there exists \tilde{Q} with $E^{\tilde{Q}}g + \theta(c, \tilde{Q}) = c < E^{\tilde{Q}}f + \theta(c, \tilde{Q})$, hence with $E^{\tilde{Q}}(g - f) < 0$. Since $\theta(c, \tilde{Q})$ is strictly binding, the criterion imposes that $\theta(0, \tilde{Q}) = 0$, and consequently $g - f \prec 0$.

20.9 Proof of Theorem 5

If part: FCC is the implication $V_1^*(f) \leq V_1(g) \Rightarrow f \preceq g$. Indeed, by TCC, $V(f) \leq V(V_1^*(f)) \leq V(V_1(g)) \leq V(g)$ when $V_1^*(f) \leq V_1(g)$. Similarly, TCC implies that $V(g) \leq V^*(V_1(g)) \leq V^*(V_1(f)) \leq V^*(f)$ when $V_1(f) \geq V_1(g)$, and BCC follows.

Only-if part: From FCC with $f = V_1(g)$ and $g \in \mathcal{A}$, it follows that $V(V_1(g)) \leq V(g)$, and from BCC we have $V(g) \leq V^*(V_1(g))$. The third inequality in TCC is equivalent to $V(V_1^*(f)) \geq V(f)$, which follows from FCC with $g = V_1^*(f)$. The last inequality in TCC is equivalent to the first (and added for symmetry).

References

- M. Allais. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica*, 21 (1953), 21:503–546, 1953.
- P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9:203–228, 1999.
- T. R. Bielecki, I. Cialenco, I. Iyigunler, and R. Rodriguez. Dynamic conic finance: Pricing and hedging in market models with transaction costs via dynamic coherent acceptability indices. *International Journal of Theoretical and Applied Finance*, 16(01):1350002, 2013.
- M. H. Birnbaum and S. E. Stegner. Source credibility in social judgment: Bias, expertise, and the judge's point of view. *Journal of Personality and Social Psychology*, 37(1):48, 1979.
- S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, and L. Montrucchio. Uncertainty averse preferences. *Journal of Economic Theory*, 146(4):1275–1330, 2011a.
- S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, and L. Montrucchio. Complete monotone quasiconcave duality. *Mathematics of Operations Research*, 36(2):321–339, 2011b.
- J. Eichberger, S. Grant, and D. Kelsey. Updating Choquet beliefs. *Journal of Mathematical Economics*, 43(7):888–899, 2007.
- D. Ellsberg. Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, pages 643–669, 1961.
- L. G. Epstein and M. Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003.
- H. Föllmer and A. Schied. *Stochastic Finance. An Introduction in Discrete Time* (3rd ed.). Walter de Gruyter, Berlin, 2011.
- M. Frittelli and M. Maggis. Dual representation of quasi-convex conditional maps. *SIAM Journal on Financial Mathematics*, 2(1):357–382, 2011.
- I. Gilboa. Rationality and the bayesian paradigm. *Journal of Economic Methodology*, 22 (3):312–334, 2015.

- I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
- E. Hanany and P. Klibanoff. Updating preferences with multiple priors. *Theoretical Economics*, 2(3):261–298, 2007.
- M. Horie. Reexamination on updating Choquet beliefs. *KIER Discussion Paper*, 643, 2007.
- 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.
- D. Kahneman, J. L. Knetsch, and R. H. Thaler. Experimental tests of the endowment effect and the coase theorem. *Journal of political Economy*, 98(6):1325–1348, 1990.
- E. Karni and Z. Safra. Preference reversals and the observability of preferences by experimental methods. *Econometrica*, pages 675–685, 1987.
- E. Karni and Z. Safra. Behaviorally consistent optimal stopping rules. *Journal of Economic Theory*, 51(2):391–402, 1990.
- M. J. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4):1622–1668, 1989.
- M. J. Machina and W. K. Viscusi. *Handbook of the Economics of Risk and Uncertainty*. Newnes, 2013.
- D. B. Madan. Relativities in financial markets. Robert H. Smith School Working Paper no. RHS-2732825, 2016.
- 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.
- C. P. Pires. A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2):137–152, 2002.
- 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.
- B. Roorda and J. M. Schumacher. Membership conditions for consistent families of non-

- etary valuations. *Statistics & Risk Modeling*, 30:255–280, 2013.
- B. Roorda and J. M. Schumacher. Weakly time consistent concave valuations and their dual representations. *Finance and Stochastics*, 20(1):123–151, 2016.
- U. Segal and A. Spivak. First order versus second order risk aversion. *Journal of Economic Theory*, 51(1):111–125, 1990.
- M. Siniscalchi. Vector expected utility and attitudes toward variation. *Econometrica*, 77(3), 2009.
- M. Siniscalchi et al. Bayesian updating for general maxmin expected utility preferences. *Unpublished paper, Department of Economics, Princeton University.[399, 400]*, 2001.
- S. Trautmann and P. Wakker. Making the Anscombe-Aumann approach to ambiguity suited for descriptive applications. Technical report, Erasmus University, 2015.
- A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.
- A. Tversky and R. H. Thaler. Anomalies: preference reversals. *The Journal of Economic Perspectives*, 4(2):201–211, 1990.
- P. P. Wakker. *Prospect theory: For risk and ambiguity*. Cambridge university press, 2010.
- P. Walley. *Statistical reasoning with imprecise probabilities*, volume 42. Peter Walley, 1991.