Membership Conditions for Consistent Families of Monetary Valuations

Berend Roorda* Hans Schumacher†

March 7 2013

Abstract

We investigate time consistency of monetary valuations, also called monetary risk measures or monetary utility functions. Through a number of recent research contributions, it has become clear that time consistency imposes strong constraints on families of monetary valuations conditioned on available information at different time instants. In this paper we add to these results by showing that consistent families of monetary valuations are already determined uniquely by the choice of the initial valuation, under suitable sensitivity assumptions; moreover, this statement holds even when the term "consistency" is interpreted in a rather weak sense. The unique update rule is specified explicitly, and we characterize the existence of consistent updates for a given initial monetary valuation. We give examples of situations in which weak consistency is relevant. An application is given to the construction of consistent families of compound valuations, as an illustration of the additional flexibility under weak time consistency.

Keywords: dynamic risk measures; nonlinear pricing; updating; weak time consistency.

1 Introduction

Risk measures and more generally valuation functionals are used for various purposes, including regulation, margin setting, asset pricing, and contract design; see for instance [2, 8, 4, 11].

*B. Roorda, School of Management and Governance, Department of Industrial Engineering and Business Information Systems, University of Twente, P.O. Box 217, 7500 AE, Enschede, the Netherlands. Phone: +3153-4894383. E-mail: b.roorda@utwente.nl. Research supported in part by Netspar.

†J.M. Schumacher, Department of Econometrics and Operations Research, CentER, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, the Netherlands. Research supported in part by Netspar.

In many applications, it is of interest to apply valuations at different points in time. Updating rules and time consistency of convex and coherent risk measures have been investigated extensively in recent years; see for instance [15, 21, 24, 12, 10, 3]. Attention has been paid in particular to the notion of *strong time consistency*, also called *dynamic consistency*. We will argue below in subsection 3.4 that, while this notion has a solid intuitive foundation in some contexts, weaker notions of consistency can still be appropriate in particular in situations in which an inequality of valuations does not necessarily imply a preference relationship.

As argued in [22], strong consistency inhibits the construction of families of risk measures that maintain comparable standards of prudence on different scales; a "VaR of VaR's", so to say, is likely to be very conservative. For related comments see also Schied [25, Rem. 3.5]. In fact, it has been noted in the literature that many risk measures do not even allow strongly consistent updates. Klöppel and Schweizer [18, Section 7.2] introduce a natural coherent risk measure based on one-sided moments, and show that it cannot be updated in a dynamically consistent way. A striking result by Kupper and Schachermayer [19] shows that, under mild technical conditions, law-invariant risk measures allow strongly time consistent updates only when they belong to the family of entropic risk measures, which is parametrized by a single scalar parameter.

Alternative, weaker notions of time consistency have been proposed and discussed in several papers, for instance [7, 22, 28, 27, 1, 5]. The main notions used in this paper are *sequential consistency* and *conditional consistency*. Both consistency notions have been introduced in [22] in the context of coherent risk measures defined on a finite outcome space. Sequential consistency is the central notion in this paper; it formalizes the idea that a position that is surely (un)acceptable at some future date should be deemed (un)acceptable already now. The importance of conditional consistency derives from the fact that, even under this weak form of consistency, updates are unique.

The uniqueness of updates, under suitable sensitivity assumptions, is one of the main results of the present paper. The result on uniqueness is supported by the construction of an operator that provides the update if it exists. This operator, called the *refinement update*, is a generalization of the well known Bayesian updating rule. We give necessary and sufficient conditions for the existence of consistent updates, and we show that consistent updating of an initial risk measure is enough to construct consistent families of risk measures. The paper concludes with an example of the construction of consistent families.

In this paper we consider the positions specified by *payoffs* (random variables) rather than by *payoff streams* (random processes) as for instance in [10] and [17]. As in most of the literature on risk measures, we shall limit ourselves to bounded random variables; methods for extending results from this case to the unbounded case are provided in [9].

The literature on risk measures that has developed following the work of Artzner et al.

[2] is marked by variations in sign conventions and terminology. The sign conventions that we use are the same as for instance in [10]: the outcomes of random variables are interpreted as gains, and positive values of functionals correspond to acceptable positions. The term "monetary utility function" that has been used in a number of recent papers is a little long-winded as noted by Jobert and Rogers [17], and moreover does not quite match the view taken in this paper that valuations are not necessarily preference indicators. We follow the suggestion in [17] to use the term "valuation", even though this term may be less suitable for applications in regulation. The term "monetary" is added as a reference to the translation axiom that is used below, and in this way we arrive at the term "monetary valuation" for the class of functionals that we consider.

The paper is organized as follows. Preliminaries with mostly well known material are presented in Section 2. The notions of consistency that we use are defined in Section 3, which we conclude by two examples that motivate the relaxation of the standard concept. The refinement update is introduced in Section 4 as the unique candidate for a consistent update. Existence of such an update is addressed in Section 5. An example of the construction of a consistent family is shown in Section 6. Finally, conclusions follow in Section 7. Most of the proofs have been collected, together with a few auxiliary results, in the Appendix.

2 Basic definitions and properties

In this section we list some basic definitions and properties and fix notation. Most of the material is standard and the basic properties are well known (see for instance [12, 10, 14]).

2.1 Standing assumptions and notation

We work in the standard setting of a fixed filtered probability space $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{t\in\mathcal{T}})$; the parameter set $\mathcal{T}$ can be an interval $[0,T]$ or a discrete set $\{t_0, t_1, \ldots, t_n\}$, with $t_0 = 0$ and $t_n = T$. We will always assume that $\mathcal{F}_0$ is trivial and $\mathcal{F}_T = \mathcal{F}$. The terms "measurable" and "almost surely" without further specification mean $\mathcal{F}$-measurable and P-almost surely, respectively. The complement of an event $F \in \mathcal{F}$ is denoted by F^c. We write $L^\infty = L^\infty(\Omega, \mathcal{F}, P)$. Elements of L^∞ will be referred to as random variables but also as "payoffs" or "positions". Convergence is taken in the almost sure sense unless indicated otherwise. All equalities and inequalities applied to random variables are understood to hold almost surely; the notation $X \lneqq Y$ means that $P(X \le Y) = 1$ and $P(X < Y) > 0$.

Given a nonempty set $\mathcal{S} \subset L^\infty$, $\operatorname{ess\,sup} \mathcal{S}$ is defined as the least element in the a.s.-equivalence classes of measurable functions from Ω to $\mathbb{R} \cup \{\infty\}$ that dominate all elements of $\mathcal{S}$ in the almost sure sense (see for instance [14]); $\operatorname{ess\,inf} \mathcal{S}$ is defined similarly. We use $\inf X$ and $\sup X$ to refer to the essential infimum and the essential supremum, respectively, of an element X of L^∞. We also use inf and sup in the usual sense to refer to the infimum and supremum of a collection of real numbers; this should not lead to confusion.

The set $L^\infty(\Omega, \mathcal{F}_t, P)$ of essentially bounded $\mathcal{F}_t$-measurable functions will be written as L^∞_t. Conditional expectations under a probability measure $Q \ll P$ are usually written as $E^Q_t X$ rather than as $E^Q[X \mid \mathcal{F}_t]$.

Given a random variable $X \in L^\infty$, the variable $\|X\|_t \in L^\infty_t$ defined by $\|X\|_t = \operatorname{ess\,inf}\{m \in L^\infty_t \mid m \geq |X|\}$ is referred to as the $\mathcal{F}_t$*-conditional norm* of X. The notation $\|X\|$ (without subscript) refers to the usual L^∞-norm of X. Since $\mathcal{F}_t \subset \mathcal{F}$, we have $L^\infty_t \subset L^\infty$ and $\|X\|_t \leq \|X\|$ for all $X \in L^\infty$.

2.2 The class $\mathcal{M}_t$ of Conditional Monetary Valuations

Definition 2.1 A *conditional monetary valuation with respect to* $\mathcal{F}_t$, also called an $\mathcal{F}_t$*-conditional monetary valuation*, is a mapping $\phi_t : L^\infty \to L^\infty_t$ that satisfies the properties of *normalization* (2.1), *monotonicity* (2.2), and $\mathcal{F}_t$*-translation invariance* (2.3):

$$\phi_t(0) = 0 \tag{2.1}$$

$$X \leq Y \Rightarrow \phi_t(X) \leq \phi_t(Y) \quad (X, Y \in L^\infty) \tag{2.2}$$

$$\phi_t(X + C_t) = \phi_t(X) + C_t \quad (X \in L^\infty,\ C_t \in L^\infty_t). \tag{2.3}$$

We use the notation $\mathcal{M}_t$ for the class of $\mathcal{F}_t$-conditional monetary valuations.

For elements of $\mathcal{M}_0$, i.e. mappings from L^∞ to $\mathbb{R}$, the term "unconditional monetary valuation" will sometimes be used. An element ϕ_t of $\mathcal{M}_t$ is called a *concave* $\mathcal{F}_t$*-conditional monetary valuation* if it satisfies $\mathcal{F}_t$*-concavity*:

$$\phi_t(\Lambda_t X + (1-\Lambda_t)Y) \geq \Lambda_t \phi_t(X) + (1-\Lambda_t)\phi_t(Y) \quad (X, Y \in L^\infty;\ \Lambda_t \in L^\infty_t,\ 0 \leq \Lambda_t \leq 1), \tag{2.4}$$

and $\phi_t \in \mathcal{M}_t$ is called *coherent* if it in addition satisfies $\mathcal{F}_t$*-positive homogeneity*:

$$\phi_t(\Lambda_t X) = \Lambda_t \phi_t(X) \quad (X \in L^\infty;\ \Lambda_t \in L^\infty_t,\ \Lambda_t \geq 0). \tag{2.5}$$

The $\mathcal{F}_t$*-local property* is always satisfied by elements of $\mathcal{M}_t$ [12, Prop. 1,2], [10, Prop. 3.3]:

$$\phi_t(1_F X + 1_{F^c} Y) = 1_F \phi_t(X) + 1_{F^c} \phi_t(Y) \quad (F \in \mathcal{F}_t;\ X, Y \in L^\infty). \tag{2.6}$$

Under the normalization assumption, the local property is equivalent to $\mathcal{F}_t$*-regularity* [12, Prop. 1]:

$$\phi_t(1_F X) = 1_F \phi_t(X) \quad (F \in \mathcal{F}_t;\ X \in L^\infty). \tag{2.7}$$

Additional assumptions relating to monotonicity that will be used frequently are *sensitivity* (2.8) and *strong sensitivity* (2.9):

$$X \lneqq 0 \ \Rightarrow\ \phi_t(X) \lneqq 0 \qquad (X \in L^\infty) \tag{2.8}$$

$$X \lneqq Y \;\Rightarrow\; \phi_t(X) \lneqq \phi_t(Y) \quad (X, Y \in L^\infty). \tag{2.9}$$

A mapping that is monotonic and strongly sensitive is said to be *strictly monotonic*. A mapping ϕ_t is *continuous from above* if

$$X_n \searrow X \Rightarrow \phi_t(X_n) \searrow \phi_t(X) \quad (X_n \in L^\infty, n = 1, 2, \ldots;\; X \in L^\infty). \tag{2.10}$$

A *dynamic monetary valuation* corresponding to the filtration $(\mathcal{F}_t)_{t\in\mathcal{T}}$ is a family $(\phi_t)_{t\in\mathcal{T}}$ with $\phi_t \in \mathcal{M}_t$ for each $t \in \mathcal{T}$. The class of dynamic monetary valuations for a given filtered probability space $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{t\in\mathcal{T}})$ is denoted by $\mathcal{M}_\mathcal{T}$.

2.3 Acceptance sets and conditional requirements

The *acceptance set* of a normalized monotonic mapping $\phi : L^\infty \to L^\infty$ is defined by

$$\mathcal{A}(\phi) = \{X \in L^\infty \mid \phi(X) \geq 0\}.$$

We will frequently write $\mathcal{A}_t$ as an abbreviation of $\mathcal{A}(\phi_t)$, when the mapping ϕ_t is clear from context. The acceptance set of an $\mathcal{F}_t$-conditional monetary valuation satisfies three properties that we express here for a general set $\mathcal{S} \subset L^\infty$, namely *acceptance of zero* (2.11), *solidness* (2.12), and *$\mathcal{F}_t$-nonnegativity* (2.13):[1]

$$0 \in \mathcal{S} \tag{2.11}$$

$$X \in \mathcal{S},\; Y \geq X \Rightarrow Y \in \mathcal{S} \quad (Y \in L^\infty) \tag{2.12}$$

$$X \in L_t^\infty \cap \mathcal{S} \Rightarrow X \geq 0. \tag{2.13}$$

Below we shall refer to these three properties as the "basic conditions". Such acceptance sets always have the *$\mathcal{F}_t$-local property* (2.14) and the *$\mathcal{F}_t$-closedness* property (2.15):

$$X, Y \in \mathcal{S} \;\Rightarrow\; 1_F X + 1_{F^c} Y \in \mathcal{S} \quad (F \in \mathcal{F}_t) \tag{2.14}$$

$$X_n \in \mathcal{S}\ (n = 1, 2, \ldots),\; \|X_n - X\|_t \to 0 \;\Rightarrow\; X \in \mathcal{S} \quad (X \in L^\infty). \tag{2.15}$$

The $\mathcal{F}_t$-closedness property follows from the inequality $|\phi_t(X) - \phi_t(Y)| \leq \|X - Y\|_t$ [10, Prop. 3.3].

The five conditions (2.11–2.15) are not only necessary but also sufficient for a set $\mathcal{S} \subset L^\infty$ to be the acceptance set of an $\mathcal{F}_t$-conditional monetary valuation. The proposition below, obtained from [12] and [10], states this fact and also explains how to relate a conditional monetary valuation to a subset $\mathcal{S}$ satisfying only the basic conditions. That construction

[1] The term "normalization" is sometimes used for properties (2.11) and (2.13) together. This phrase may be too simple however since it does not indicate that the defined notion depends on the σ-algebra $\mathcal{F}_t$.

relies on the notion of a *conditional capital requirement*, introduced in [12], which associates to an arbitrary set $\mathcal{S} \subset L^\infty$ the mapping from L^∞ to L_t^∞ given by

$$\phi_{\mathcal{S}}^t(X) = \operatorname{ess\,sup}\{Y_t \in L_t^\infty \mid X - Y_t \in \mathcal{S}\}. \tag{2.16}$$

For the convenience of the reader, the proof of the proposition is summarized in the Appendix.

Proposition 2.2 *A set $\mathcal{S} \subset L^\infty$ is the acceptance set of an $\mathcal{F}_t$-conditional monetary valuation if and only if it satisfies the five conditions (2.11–2.15). The associated conditional monetary valuation is uniquely determined as the capital requirement $\phi_{\mathcal{S}}^t$ of $\mathcal{S}$, defined by (2.16). More generally, for any $\mathcal{S} \subset L^\infty$ satisfying the basic conditions (2.11–2.13), $\phi_{\mathcal{S}}^t$ is the element of $\mathcal{M}_t$ whose acceptance set is equal to the smallest extension of $\mathcal{S}$ that satisfies (2.14–2.15). If in addition $\mathcal{S}$ is convex , then $\phi_{\mathcal{S}}^t$ is a concave valuation.*

It follows that the construction of capital requirements induces a one-to-one correspondence between conditional monetary valuations $\phi_t \in \mathcal{M}_t$ and their acceptance sets, given by

$$\phi_t = \phi_{\mathcal{A}(\phi_t)}^t. \tag{2.17}$$

The acceptance set of a concave valuation is convex, and in fact satisfies the stronger property of $\mathcal{F}_t$-*convexity* which is expressed as follows:

$$\Lambda_t X + (1 - \Lambda_t) Y \in \mathcal{S} \quad (X, Y \in \mathcal{S};\ \Lambda_t \in L_t^\infty,\ 0 \le \Lambda_t \le 1). \tag{2.18}$$

Sensitivity of a conditional monetary valuation ϕ_t is reflected by the property of *negative cone exclusion*:

$$X \lneqq 0 \ \Rightarrow X \notin \mathcal{S}. \tag{2.19}$$

For later reference, we identify two properties that represent distinct features of the $\mathcal{F}_t$-local property (2.14). The first property, *closedness under $\mathcal{F}_t$-isolation*, is related to *restricting* a given position, while the second, *$\mathcal{F}_t$-complementarity*, relates to *joining* two mutually exclusive positions. These two properties are expressed as follows:

$$X \in \mathcal{S} \Rightarrow 1_F X \in \mathcal{S} \quad (X \in L^\infty,\ F \in \mathcal{F}_t) \tag{2.20}$$

$$1_F X \in \mathcal{S},\ 1_{F^c} X \in \mathcal{S} \Rightarrow X \in \mathcal{S} \quad (X \in L^\infty,\ F \in \mathcal{F}_t). \tag{2.21}$$

Proposition 2.3 *Let $\mathcal{F}_t$ be a sub-σ-algebra. A set $\mathcal{S} \subset L^\infty$ that satisfies $0 \in \mathcal{S}$ has the $\mathcal{F}_t$-local property if and only if it has both the $\mathcal{F}_t$-complementarity property and the property of closedness under $\mathcal{F}_t$-isolation.*

PROOF First, assume that $\mathcal{S}$ has the local property. For any $X \in \mathcal{S}$ and $F \in \mathcal{F}_t$, we have $1_F X = 1_F X + 1_{F^c} 0 \in \mathcal{S}$, so that $\mathcal{S}$ is closed under $\mathcal{F}_t$-isolation. To prove the complementarity property, let $X \in L^\infty$ and $F \in \mathcal{F}_t$ be such that $1_F X \in \mathcal{S}$ and $1_{F^c} X \in \mathcal{S}$. Writing $X = 1_F(1_F X) + 1_{F^c}(1_{F^c} X)$, we see that the local property implies that $X \in \mathcal{S}$.

Conversely, assume now that $\mathcal{S}$ is closed under $\mathcal{F}_t$-isolation and has the $\mathcal{F}_t$-complementarity property. Take $X, Y \in \mathcal{S}$, and $F \in \mathcal{F}_t$, and write $Z = 1_F X + 1_{F^c} Y \in \mathcal{S}$. We need to prove that $Z \in \mathcal{S}$. Note that $1_F Z = 1_F X \in \mathcal{S}$ and $1_{F^c} Z = 1_{F^c} Y \in \mathcal{S}$ by the closedness under $\mathcal{F}_t$-isolation of $\mathcal{S}$. By the $\mathcal{F}_t$-complementarity, this suffices to show that indeed $Z \in \mathcal{S}$. □

3 Time consistency

3.1 Sequential consistency

Several notions of time consistency are used in the literature. The notion of *sequential consistency* is central in this paper. This notion is defined as follows.

DEFINITION 3.1 Let be given $\phi_s \in \mathcal{M}_s$ and $\phi_t \in \mathcal{M}_t$ with $s \le t$. We say that ϕ_s and ϕ_t are *sequentially consistent*, or that ϕ_t is a *sequentially consistent $\mathcal{F}_t$-update* of ϕ_s, if the following conditions hold:

$$\phi_t(X) \ge 0 \;\Rightarrow\; \phi_s(X) \ge 0 \quad (X \in L^\infty) \tag{3.1a}$$

$$\phi_t(X) \le 0 \;\Rightarrow\; \phi_s(X) \le 0 \quad (X \in L^\infty). \tag{3.1b}$$

The condition (3.1a) is known as "weak acceptance consistency" [7, 13, 26] while the property (3.1b) has been called "weak rejection consistency" [26]. We shall use the simpler terms *acceptance consistency* and *rejection consistency* instead. The combination of the two properties, which we refer to as sequential consistency, was used by Weber in a study of distribution-invariant risk measures [28].

The following characterizations of sequential consistency may aid the intuition (cf. [22, Thm. 4.2], [26, Kor. 3.1.8]). Recall that we use $\inf X$ ($\sup X$) to denote the essential infimum (supremum) of an element of L^∞; in particular, $\inf X$ and $\sup X$ are constants.

LEMMA 3.2 *The monetary valuation $\phi_t \in \mathcal{M}_t$ is a sequentially consistent update of $\phi_s \in \mathcal{M}_s$ if and only if the following equivalent conditions hold:*

(i) $\phi_t(X) = 0 \Rightarrow \phi_s(X) = 0 \quad (X \in L^\infty)$

(ii) $\phi_s(X - \phi_t(X)) = 0 \quad (X \in L^\infty)$

(iii) $\inf \phi_t(X) \le \phi_s(X) \le \sup \phi_t(X) \quad (X \in L^\infty)$.

Proof Clearly, property (i) is implied by sequential consistency. For any $X \in L^\infty$ we have $\phi_t(X - \phi_t(X)) = 0$, so that property (ii) is implied by property (i). If property (ii) holds, then for any $X \in L^\infty$ we have

$$\phi_s(X) - \inf \phi_t(X) = \phi_s(X - \inf \phi_t(X)) \geq \phi_s(X - \phi_t(X)) = 0$$

and likewise $\phi_s(X) - \sup \phi_t(X) \leq 0$, so that (iii) is satisfied. Finally, it is immediate that property (iii) implies sequential consistency. □

The notion of *strict sequential consistency* is defined as sequential consistency with the added requirement

$$\phi_t(X) \gneqq 0 \Rightarrow \phi_s(X) \gneqq 0 \quad (X \in L^\infty). \tag{3.2}$$

It is easily verified that under strong sensitivity this is equivalent to sequential consistency.

3.2 Strong time consistency

The notion of time consistency that is used most frequently in the literature is *strong time consistency*, also called *dynamic consistency* or just *time consistency*; see for instance [3, Def. 5.2], [16, Def. 18], [13, Def. 3.1]. We will also refer to it as *strong time consistency.*

Definition 3.3 Let be given $\phi_s \in \mathcal{M}_s$ and $\phi_t \in \mathcal{M}_t$ with $s \leq t$. We say that ϕ_s and ϕ_t are *strongly time consistent*, or that ϕ_t is a *strongly time consistent update* of ϕ_s, if the following relation holds for all $X \in L^\infty$:

$$\phi_s(\phi_t(X)) = \phi_s(X). \tag{3.3}$$

A characterization in terms of acceptance sets is given in [13]. The definition of strong time consistency is sometimes given in the form of one of the following equivalent implications:

$$\phi_t(X) = \phi_t(Y) \Rightarrow \phi_s(X) = \phi_s(Y) \quad (X, Y \in L^\infty,\ s \leq t), \tag{3.4}$$

$$\phi_t(X) \leq \phi_t(Y) \Rightarrow \phi_s(X) \leq \phi_s(Y) \quad (X, Y \in L^\infty,\ s \leq t). \tag{3.5}$$

Under $\mathcal{F}_t$-translation invariance, these conditions are both equivalent to the definition above, as can be seen by taking $Y = \phi_t(X)$, cf. [1, Prop 1.16]. It is immediately clear from the definitions that strong time consistency implies sequential consistency. Reasons for considering the weaker notion will be discussed in subsection 3.4 below.

3.3 Conditional consistency

Finally we introduce a notion that is even weaker (under suitable sensitivity assumptions) than sequential consistency. This notion plays a key role in uniqueness of updating.

Definition 3.4 Let be given $\phi_s \in \mathcal{M}_s$ and $\phi_t \in \mathcal{M}_t$ with $s \le t$. We say that ϕ_s and ϕ_t are *conditionally consistent*, or that ϕ_t is a *conditionally consistent $\mathcal{F}_t$-update* of ϕ_s, if the following condition holds:

$$\phi_t(X) \ge 0 \;\Leftrightarrow\; \forall F \in \mathcal{F}_t : \phi_s(1_F X) \ge 0 \quad (X \in L^\infty). \tag{3.6}$$

The condition in the definition states that approval of a position at level t is equivalent to approval at the aggregate level s not only of the position itself, but also of its isolated versions where isolation is taken up to level t.

In order to describe the notion of conditional consistency in terms of acceptance sets, we introduce the following construction.

Definition 3.5 Given a set $\mathcal{S} \subset L^\infty$ such that $0 \in \mathcal{S}$, the *$\mathcal{F}_t$-restriction* of $\mathcal{S}$ is the set $\mathcal{S}^t$ defined by

$$\mathcal{S}^t = \{X \in \mathcal{S} \,|\, 1_F X \in \mathcal{S} \text{ for all } F \in \mathcal{F}_t\}. \tag{3.7}$$

The set $\mathcal{S}^t$ can be described alternatively as the largest subset of $\mathcal{S}$ that is closed under $\mathcal{F}_t$-isolation. Definition 3.5 has been used before by Tutsch [26, p. 88], in the situation in which the set $\mathcal{S}$ is the acceptance set $\mathcal{A}_s$ of a monetary valuation $\phi_s \in \mathcal{M}_s$. She refers to this set, which we denote by $\mathcal{A}_s^t$, as the acceptance set of ϕ_s with respect to $\mathcal{F}_t$.

Conditional consistency can now be formulated compactly as the requirement that

$$\mathcal{A}_t = \mathcal{A}_s^t. \tag{3.8}$$

Since the acceptance set $\mathcal{A}_t$ is closed under $\mathcal{F}_t$-isolation as a consequence of the local property, we have $\mathcal{A}_t \subset \mathcal{A}_s^t$ if and only if $\mathcal{A}_t \subset \mathcal{A}_s$. In other words, acceptance consistency is equivalent to inclusion from left to right in (3.8). The reverse inclusion, which is called the *consecutivity property* in [26], is not equivalent to rejection consistency, however. The relations between various notions of consistency are indicated in the following proposition, whose proof is in the Appendix.

Proposition 3.6 *Let be given $\phi_s \in \mathcal{M}_s$ and $\phi_t \in \mathcal{M}_t$ with $s \le t$. Conditional consistency of ϕ_s and ϕ_t is implied in each of the following cases:*

1. *ϕ_s and ϕ_t are strongly time consistent, and ϕ_s is sensitive*
2. *ϕ_s and ϕ_t are sequentially consistent, and ϕ_s is strongly sensitive*
3. *ϕ_s and ϕ_t are strictly sequentially consistent.*

The following proposition shows that conditional consistency is strong enough to preserve some properties of interest.

Proposition 3.7 *Let be given a sensitive monetary valuation $\phi_s \in \mathcal{M}_s$, and let ϕ_t be a conditionally consistent $\mathcal{F}_t$-update of ϕ_s. Then the following statements hold:*

(i) ϕ_t is sensitive;

(ii) if ϕ_s is concave, then so is ϕ_t;

(iii) if ϕ_s is continuous from above, then so is ϕ_t.

The proof is given in the Appendix. It is also shown in the Appendix (Lemma 8.5) that strong sensitivity is preserved under sequentially consistent updating.

3.4 Strong versus weak time consistency

The characterizations (3.4) and (3.5) of strong time consistency are often seen as having a solid intuitive foundation. In [9], for instance, it is argued that "a violation of this condition clearly leads to capital requirements that are inconsistent over time". Also in [12] it is concluded that "the financial meaning of time consistency is based on general intuition". However, different contexts may lead to different intuitions. We present two examples below in which violations of (3.5) occur without apparent contradictions at least in our view.

Example 3.8 Consider a nonrecombining two-step binomial tree, with 1% probability for all downward branches. Let $\hat{\phi}_0$ and $\hat{\phi}_1$ denote the single-step worst-case operators between time 0 to time 1 and between time 1 to time 2, respectively. Define a conditional monetary valuation ϕ by

$$\phi_1 = \hat{\phi}_1, \qquad \phi_0(X) = \min\big(E_0[\hat{\phi}_1(X)],\, \hat{\phi}_0(E_1[X])\big) \quad (X \in L^\infty) \tag{3.9}$$

(cf. Sequential TVaR as proposed in [22]). This conditional valuation is sequentially consistent but not strongly time consistent. Indeed, compare the two positions (in obvious notation, from up-up to down-down) $X = (0, 0, 0, -1)$ and $Y = (0.1, 0.1, -0.9, -0.9)$. We have $\phi_1(X) = (0, -1) < \phi_1(Y) = (0.1, -0.9)$ whereas $\phi_0(X) = -0.01 > \phi_0(Y) = -0.9$. Nevertheless, from a risk management perspective at 99% confidence level, these outcomes do not seem unreasonable. It is important to note that the inequality $\phi_0(X) > \phi_0(Y)$ does not necessarily mean that an entity that holds position Y at time 0 would be happy to replace Y by X; other criteria besides the one given by the functional ϕ_0 may also play a role in determining the attractiveness of positions.

We may compare the outcomes to the ones obtained from the strongly consistent conditional valuation that is based on the single-step operators $\hat{\phi}_0$ and $\hat{\phi}_1$:

$$\phi'_1 = \hat{\phi}_1, \qquad \phi'_0(X) = \hat{\phi}_0(\hat{\phi}_1(X)) \quad (X \in L^\infty). \tag{3.10}$$

We have $\phi'_0(X) = -1 < \phi'_0(Y) = -0.9$. The evaluation of the position X appears very conservative.

The substantially different positions $(0, 0, -1, -1)$ and $(0, 0, 0, -1)$ are equally evaluated under ϕ'_0. Under strong consistency it has to be this way, because these two positions are equally evaluated under ϕ'_1. In other words, strong consistency implies no distinction can be made at an anterior node between on the one hand a *reserve* that should be maintained at a given later node to cover a potential loss, and on the other hand a *sure loss* of the same size at that node.

EXAMPLE 3.9 In this example we think of valuation functionals as producing ask prices for financial contracts. Consider a stochastic payoff X in the (distant) future; let $\phi_0(X)$ be its ask price at time 0, and let $\phi_1(X)$ denote its ask price at time 1. We can also consider the ask price at time 0 of the contract that pays $\phi_1(X)$ at time 1. This contract allows the holder to buy at time 1 the asset that pays X at time T, but the holder may as well use the payoff in a different way. If the asset paying X was already bought at time 0, conversion to cash at time 1 would only be possible at the *bid* price of the asset at time 1, which may be considerably lower than the ask price. It would be reasonable therefore that the inequality

$$\phi_0(\phi_1(X)) > \phi_0(X)$$

holds. Take $\delta > 0$ such that $\phi_0(\phi_1(X)) - \delta > \phi_0(X)$, and consider the contract Y that pays $\phi_1(X) - \delta$ at time 1. Obviously we have $\phi_1(Y) < \phi_1(X)$, while at the same time $\phi_0(Y) = \phi_0(\phi_1(X)) - \delta > \phi_0(X)$. There appears to be no contradiction though. If the *bid* price of X at time 1 would be higher than the cash amount Y, then certainly X would be preferred to Y at that time, but the inequality is in terms of *ask* prices. As in the previous example, it seems to be essential that inequality of valuations does not necessarily imply a preference relation.

4 Uniqueness of updating

It is a well known fact, recalled in Prop. 2.2 above, that any given conditional monetary valuation can be viewed as the conditional capital requirement corresponding to its acceptance set. The construction of the $\mathcal{F}_t$-restriction, introduced in Section 3 to express conditional consistency in terms of acceptance sets, modifies a given acceptance set in a way that relates to the filtration member $\mathcal{F}_t$. This suggests a particular way of updating a given monetary valuation. Namely, given $\phi_s \in \mathcal{M}_s$, take its acceptance set $\mathcal{A}_s$, construct the $\mathcal{F}_t$-restriction $\mathcal{A}_s^t$, and define the $\mathcal{F}_t$-update of ϕ_s as the conditional capital requirement that corresponds to $\mathcal{A}_s^t$. In fact, if the update is to be conditionally consistent, this construction provides the *only* feasible candidate, since by definition of conditional consistency the acceptance set of the update must be equal to the $\mathcal{F}_t$-restriction of the acceptance set of the monetary valuation at the earlier time. Since as we have seen above conditional consistency is implied, under suitable sensitivity assumptions, by strong time consistency as well as by sequential consistency, this implies that we obtain a criterion for a given monetary valuation to be a member of a sequentially or strongly time consistent family of monetary valuations. Namely, it suffices to check whether the update that was described above has the property.

To carry out this program, we first of all need to check whether the set $\mathcal{A}_s^t$ does indeed define a conditional capital requirement. Conditions for this to be the case have been formulated in Prop. 2.2. We verify that these conditions are satisfied.

Proposition 4.1 *If $\mathcal{S} \subset L^\infty$ satisfies the three properties (2.11) (acceptance of zero), (2.12) (solidness), and (2.19) (negative cone exclusion), then the $\mathcal{F}_t$-restriction $\mathcal{S}^t$ of $\mathcal{S}$ satisfies (2.11) and (2.12) as well, and moreover S^t has the conditional nonnegativity property (2.13) with respect to $\mathcal{F}_t$.*

Proof The inheritance of the properties of solidness and acceptance of zero is trivial. To show the conditional nonnegativity property, suppose there exists $X_t \in L_t^\infty \cap \mathcal{S}^t$ such that $X_t \not\geq 0$. Then there exist $\varepsilon > 0$ and $F \in \mathcal{F}_t$ with $P(F) > 0$ such that $1_F X_t \leq -\varepsilon 1_F$. Since $X_t \in \mathcal{S}^t$ and $F \in \mathcal{F}_t$, we have $1_F X_t \in \mathcal{S}$. By the solidness of $\mathcal{S}$ it then follows that $-\varepsilon 1_F \in \mathcal{S}$, which is incompatible with the negative cone exclusion property (2.19). □

It is easily verified that the proposition is applicable to the acceptance set $\mathcal{A}_s$ of a sensitive $\mathcal{F}_s$-conditional monetary valuation.

Corollary 4.2 *If $\phi_s \in \mathcal{M}_s$ is sensitive, then the $\mathcal{F}_t$-restriction $\mathcal{A}_s^t$ of its acceptance set $\mathcal{A}_s$ satisfies the basic properties of Prop. 2.2, namely acceptance of zero, solidness, and conditional nonnegativity.*

As a consequence, the following definition is justified.

Definition 4.3 Let a sensitive $\mathcal{F}_s$-conditional monetary valuation ϕ_s be given, and let $t \geq s$. The *$\mathcal{F}_t$-refinement update* of ϕ_s is the $\mathcal{F}_t$-conditional monetary valuation ϕ_s^t defined by

$$\phi_s^t(X) = \operatorname{ess\,sup}\{Y \in L_t^\infty \mid \phi_s(1_F(X - Y)) \geq 0 \text{ for all } F \in \mathcal{F}_t\}. \tag{4.1}$$

The uniqueness of updating that was already mentioned above can now be stated more formally. We know, as recalled in Prop. 2.2, that a monetary valuation is the conditional capital requirement of its acceptance set, so the refinement update provides the unique candidate for a conditionally consistent update, having $\mathcal{A}_t = \mathcal{A}_s^t$. In view of Prop. 3.6, this must also hold for (strict) sequentially and strongly time consistent updates, under appropriate sensitivity conditions.

THEOREM 4.4 *Let a sensitive $\mathcal{F}_s$-conditional monetary valuation ϕ_s be given, and let $t \geq s$. The refinement update ϕ_s^t of ϕ_s is the unique conditionally, strictly sequentially, and/or strongly time consistent $\mathcal{F}_t$-update of ϕ_s, if such an update of ϕ_s exists. If ϕ_s is strongly sensitive, this claim also holds for the sequentially consistent $\mathcal{F}_t$-update of ϕ_s.*

Uniqueness of *strongly* time consistent updating has been proved in a different way by Cheridito et al. [10, Cor. 4.8]. Uniqueness of *sequentially* consistent updating was proved under some technical conditions for distribution-invariant risk measures by Weber [28, Cor. 4.1].

Application of the theorem in a particular case may or may not be straightforward, depending on whether the refinement update is easily computed and the desired type of consistency can be easily verified. As an illustration, we give an example of a valuation that does not allow a conditionally consistent update.

EXAMPLE 4.5 Consider a nonrecombining two-step binomial tree with probability $\frac{1}{2}$ for all branches; we use notation as in Example 3.8. Define an unconditional concave valuation $\phi_0 : \mathbb{R}^4 \to \mathbb{R}$ by

$$\phi_0(X) = \min(\tfrac{1}{4}(x_1 + x_2 + x_3 + x_4), \tfrac{1}{6}(x_1 + 2x_2 + 2x_3 + x_4 + 1)).$$

This valuation is strongly sensitive. Calculation shows that the refinement update $\phi_1 : \mathbb{R}^4 \to \mathbb{R}^2$ is given by

$$\phi_1(X) = \big(\min(\tfrac{1}{2}(x_1 + x_2), \tfrac{1}{3}(x_1 + 2x_2 + 1)), \min(\tfrac{1}{2}(x_3 + x_4), \tfrac{1}{3}(2x_3 + x_4 + 1))\big).$$

To see that the refinement update is not conditionally consistent, take $X = (1, -1, -1, 1)$; we have $\phi_1(X) = (0, 0)$ while $\phi_0(X) = -1$. In other words, even acceptance consistency does not hold. Therefore, the valuation ϕ_0 does not admit a conditionally consistent update, and neither does it allow a sequentially consistent or a strongly time consistent update. Looking at the situation in more detail, we can see that both vectors $(1, -1, 0, 0)$ and $(0, 0, -1, 1)$ belong to $\mathcal{A}_0^1$, but their sum does not. Therefore $\mathcal{A}_0^1$ does not satisfy the $\mathcal{F}_1$-complementarity property, and hence cannot be the acceptance set of an$\mathcal{F}_1$-conditional monetary valuation.

Consider now a dynamic monetary valuation $\phi \in \mathcal{M}_{\mathcal{T}}$ associated to the given filtration $(\mathcal{F}_t)_{t\in\mathcal{T}}$, that is, a family $(\phi_t)_{t\in\mathcal{T}}$ such that, for each $t \in \mathcal{T}$, $\phi_t \in \mathcal{M}_t$. It would be reasonable to speak of a *strongly time consistent dynamic valuation* if, for each $s, t \in \mathcal{T}$ with $s \leq t$, ϕ_t is a strongly time consistent update of ϕ_s, and likewise for other notions of consistency. The conditions in Thm. 4.4 provide criteria for an initial valuation ϕ_0 to admit consistent updates with respect to more detailed σ-algebras $(\mathcal{F}_t)_{t\in\mathcal{T}}$, but it remains to be seen whether in this way a consistent family is defined, since in principle it may happen that ϕ_0 is consistent with ϕ_s and also with ϕ_t, but ϕ_s and ϕ_t are not consistent with each other. In fact simple examples show that such situations may indeed arise if the notion of consistency that is used is acceptance consistency. The same holds if one uses the notion of middle rejection consistency which is defined [20, Def. 2.1.2, Prop. 2.1.6] (cf. also [26, Thm. 3.1.5]) by the condition $\phi_s(X) \leq \phi_s(\phi_t(X))$ for $X \in L^\infty$ and $s \leq t$. In the case of conditional consistency, however, consistent updating of ϕ_0 is enough to construct a consistent family. Under strong sensitivity, the same holds for sequential consistency and for strong time consistency. This is a consequence of the following proposition.

Proposition 4.6 *Let be given $\phi_s \in \mathcal{M}_s$, $\phi_t \in \mathcal{M}_t$, and $\phi_u \in \mathcal{M}_u$, with $s \leq t \leq u$. If ϕ_t is a conditionally consistent update of ϕ_s, then ϕ_u is a conditionally consistent update of ϕ_t if and only if it is a conditionally consistent update of ϕ_s. Under the assumption that ϕ_s is strongly sensitive, the same statement holds when "conditionally consistent" is replaced throughout either by "(strictly) sequentially consistent" or by "strongly time consistent".*

The proof of the proposition is given in the Appendix.

5 Membership of consistent families

It was shown above that the refinement update provides the unique candidate for consistent updating. As shown in Example 4.5, the refinement update need not be conditionally consistent, let alone sequentially or strongly time consistent. This means that, given an unconditional valuation $\phi_0 \in \mathcal{M}_0$, it is not in general true that this valuation is a member of a consistent family of valuations, even if we require only conditional consistency. On the other hand, as shown in the previous section, membership of such a family does follow if we can show that consistent updating of the initial valuation ϕ_0 is possible with respect to the σ-algebras belonging to the given filtration $(\mathcal{F}_t)_{t\in\mathcal{T}}$.

Given the uniqueness result Thm. 4.4, to decide whether or not a conditionally consistent update exists it is sufficient to compute the refinement update and to check whether it satisfies the requirements (3.6). In fact it is enough to verify acceptance consistency, as shown in the following proposition.

Proposition 5.1 *Let ϕ_t be the $\mathcal{F}_t$-refinement update of a given $\phi_s \in \mathcal{M}_s$. Then ϕ_t is a conditionally consistent update of ϕ_s if it is an acceptance consistent update.*

Proof Suppose ϕ_t is the $\mathcal{F}_t$-refinement update of ϕ_s and (3.1a) holds. The implication from right to left in (3.6) follows from the definition of the refinement update. Moreover (as noted by Tutsch [26, Kor. 3.1.8(d')]), if $\phi_t(X) \geq 0$, then for any $F \in \mathcal{F}_t$ also $\phi_t(1_F X) = 1_F \phi_t(X) \geq 0$, and we can conclude that $\phi_s(1_F X) \geq 0$ by applying (3.1a) to $1_F X$. $\square$

In other words, the proposition states that if the refinement update of a given monetary valuation is not conditionally consistent, then it is in fact not even acceptance consistent. This is the situation we encountered in Example 4.5. In [26, Remark 8.2] a similar observation has been made for the class of convex risk measures.

In cases in which the computation of the refinement update and verification of acceptance consistency is not easily achieved, it is of interest to have alternative conditions for the existence of conditionally consistent updates. According to Prop. 2.2, there are five properties that need to be satisfied in order for the $\mathcal{F}_t$-restriction of the acceptance set $\mathcal{A}_s$ to qualify as the acceptance set of an $\mathcal{F}_t$-conditional monetary valuation. The three "basic conditions" mentioned at the beginning of subsection 2.3 are always satisfied for an $\mathcal{F}_t$-update of an acceptance set. The two remaining properties are $\mathcal{F}_t$-closedness and the $\mathcal{F}_t$-local property. As may be expected, the closedness property holds under a continuity assumption. The proof of the lemma below can be found in the Appendix.

Lemma 5.2 *If $\phi : L^\infty \to L^\infty$ is normalized, monotonic, and continuous from above, then, for any σ-algebra $\mathcal{F}_t \subset \mathcal{F}$, the $\mathcal{F}_t$-restriction of the acceptance set of ϕ is $\mathcal{F}_t$-closed.*

The existence of conditional updates can therefore be characterized as follows for valuations that are continuous from above.

Proposition 5.3 *If $\phi_s \in \mathcal{M}_s$ is sensitive and continuous from above, then ϕ_s admits a conditionally consistent $\mathcal{F}_t$-update if and only if it has the following property:*

$$[\,\exists G \in \mathcal{F}_t : \; \forall F \in \mathcal{F}_t : \; \phi_s(1_{F\cap G}X) \geq 0, \; \phi_s(1_{F\cap G^c}X) \geq 0\,] \;\Rightarrow\; \phi_s(X) \geq 0 \quad (X \in L^\infty). \tag{5.1}$$

Proof As already noted above, under the stated conditions the valuation ϕ_s admits a conditionally consistent $\mathcal{F}_t$-update if and only if the $\mathcal{F}_t$-restriction $\mathcal{A}_s^t$ of the acceptance set of ϕ_s has the $\mathcal{F}_t$-local property (2.14). It was shown in Prop. 2.3 that the $\mathcal{F}_t$-local property is equivalent to the combination of the $\mathcal{F}_t$-complementarity property (2.21) and closedness under $\mathcal{F}_t$-isolation (2.20). The latter property is always satisfied by $\mathcal{F}_t$-restrictions, so that it only remains to show that the property (5.1) is equivalent to $\mathcal{F}_t$-complementarity of $\mathcal{A}_s^t$. By definition, $\mathcal{F}_t$-complementarity of $\mathcal{A}_s^t$ means that, for $X \in L^\infty$, we have

$$[\,\exists G \in \mathcal{F}_t : \; \forall F \in \mathcal{F}_t : \; \phi_s(1_{F\cap G}X) \geq 0, \; \phi_s(1_{F\cap G^c}X) \geq 0\,] \;\Rightarrow\; \forall F \in \mathcal{F}_t : \; \phi_s(1_F X) \geq 0. \tag{5.2}$$

This property obviously implies (5.1). Conversely, suppose that (5.1) holds, and let $X \in L^\infty$ be such that the condition on the left hand side of (5.2) is fulfilled. Take $H \in \mathcal{F}_t$. The condition in (5.2) is then fulfilled also for $1_H X$ instead of X, and since this premise is the same as in (5.1) it follows that $\phi_s(1_H X) \geq 0$. Therefore the conclusion of (5.2) holds. Consequently (5.2) is equivalent to (5.1). □

We will refer to (5.1) as the property of *complementary acceptance*. This criterion is considerably weaker than the $\mathcal{F}_t$-complementarity property of the acceptance set of ϕ_s, which may be formulated as

$$[\,\exists G \in \mathcal{F}_t : \phi_s(1_G X) \geq 0,\ \phi_s(1_{G^c} X) \geq 0\,] \;\Rightarrow\; \phi_s(X) \geq 0 \quad (X \in L^\infty). \tag{5.3}$$

Indeed, the premise of (5.1) implies the one in (5.3).

If we require that the local property of the refinement should hold generically, that is, with respect to *any* filtration, then we can formulate an even simpler necessary and sufficient condition. This is a consequence of the following lemma. The property (5.4) that is used below may be called *generic complementarity.*

LEMMA 5.4 *Let $\mathcal{S} \subset L^\infty$ be such that $0 \in \mathcal{S}$. The $\mathcal{F}_t$-refinement $\mathcal{S}^t$ satisfies the $\mathcal{F}_t$-local property for all σ-algebras $\mathcal{F}_t \subset \mathcal{F}$ if and only if the following property holds:*

$$X, Y \in \mathcal{S},\ P\big(\{X = 0\} \cup \{Y = 0\}\big) = 1 \;\Rightarrow\; X + Y \in \mathcal{S}. \tag{5.4}$$

PROOF The sufficiency part follows from the discussion above. For the necessity, let $X, Y \in \mathcal{S}$ be such that $P\big(\{X = 0\} \cup \{Y = 0\}\big) = 1$. Define $F = \{X \neq 0\} \in \mathcal{F}$, and take $\mathcal{F}_t = \{\emptyset, F, F^c, \Omega\}$. We have $1_F X = X \in \mathcal{S}$ and $1_{F^c} X = 0 \in \mathcal{S}$, so that $X \in \mathcal{S}^t$, and likewise $Y \in \mathcal{S}^t$. By assumption the set $\mathcal{S}^t$ has the $\mathcal{F}_t$-local property, so that $1_F X + 1_{F^c} Y \in \mathcal{S}^t$. Since $1_F X = X$, $1_{F^c} Y = Y$, and $\mathcal{S}^t \subset \mathcal{S}$, this implies $X + Y \in \mathcal{S}$. □

A result showing that coherent risk measures that are sensitive and continuous from above can always be updated in a conditionally consistent way was proved in the context of tree models in [22, Thm. 7.1]. A more general statement follows below. The update rule that is used for this class is known as *Bayesian updating.* The proposition below therefore shows that the refinement update can be viewed as a generalization of Bayesian updating to valuations that are not of the coherent ("multiple-prior") type.

PROPOSITION 5.5 *Let $\mathcal{Q}$ be a collection of probability measures that are all absolutely continuous with respect to the reference measure P. Define*

$$\phi_t(X) = \operatorname*{ess\,inf}_{Q \in \mathcal{Q}} E_t^Q X \qquad (X \in L^\infty,\ t \in \mathcal{T}). \tag{5.5}$$

Assume that ϕ_0 is sensitive. The dynamic valuation $(\phi_t)_{t \in \mathcal{T}}$ that is defined in this way is conditionally consistent.

Proof Take $s, t \in \mathcal{T}$ with $s \leq t$. We need to show that $\phi_t(X) \geq 0$ if and only if $\phi_s(1_F X) \geq 0$ for all $F \in \mathcal{F}_t$. First, let $X \in L^\infty$ be such that $\phi_t(X) \geq 0$. Take $F \in \mathcal{F}_t$ and $Q \in \mathcal{Q}$. It follows from $\phi_t(X) \geq 0$ that $E_t^Q X \geq 0$ and hence $1_F E_t^Q X \geq 0$, so that $E_s^Q(1_F X) = E_s^Q E_t^Q(1_F X) = E_s^Q(1_F E_t^Q X) \geq 0$. Therefore we have $\phi_s(1_F X) \geq 0$ for all $F \in \mathcal{F}_t$, as required. For the converse, note that since $E_s^Q(1_F X) = E_s^Q(1_F E_t^Q X)$ for $F \in \mathcal{F}_t$, it follows from $\phi_s(1_F X) \geq 0$ for all $F \in \mathcal{F}_t$ that $\phi_s(1_F E_t^Q X)$ for all $F \in \mathcal{F}_t$. Using Lemma 8.4 in the Appendix, we can conclude from this that $E_t^Q X \geq 0$. □

The sensitivity requirement in the proposition holds if the collection $\mathcal{Q}$ contains at least one measure that is equivalent to P, and requires that the convex hull of $\mathcal{Q}$ contains such an element, cf. [13, Cor 3.6].

Necessary and sufficient conditions for the existence of sequentially consistent updates were given for the class of law-invariant concave valuations by Weber [28, Thm. 4.3, 4.4]. The strongly time consistent families within the same class have been fully described by Kupper and Schachermayer [19]. Alternative characterizations (not requiring law invariance) are provided in the following two propositions.

Proposition 5.6 *A strongly sensitive $\mathcal{F}_s$-conditional monetary valuation ϕ_s admits a sequentially consistent $\mathcal{F}_t$-update if and only if it admits a conditionally consistent update and for each $X \in L^\infty$ there exists $C_t \in L_t^\infty$ such that*

$$\phi_s(1_F(X - C_t)) = 0 \quad (F \in \mathcal{F}_t). \tag{5.6}$$

Proposition 5.7 *A strongly sensitive $\mathcal{F}_s$-conditional monetary valuation ϕ_s admits a strongly time consistent $\mathcal{F}_t$-update if and only if for each $X \in L^\infty$ there exists $C_t \in L_t^\infty$ such that*

$$\phi_s(1_F X) = \phi_s(1_F C_t) \quad (F \in \mathcal{F}_t). \tag{5.7}$$

The proofs of both propositions are in the Appendix.

6 Dynamic compound valuations

To illustrate the construction of consistent families, we discuss in this section dynamic valuations that are constructed by means of compounding. In general, starting with some collection Φ of dynamic valuations $\phi = (\phi_t)_{t\in\mathcal{T}}$, one can define a new dynamic valuation $\hat{\phi} = (\hat{\phi}_t)_{t\in\mathcal{T}}$ by

$$\hat{\phi}_t(X) = \operatorname*{ess\,inf}_{\phi\in\Phi} \phi_t(X) \qquad (X \in L^\infty). \tag{6.1}$$

It is easily verified that indeed $\hat{\phi}_t \in \mathcal{M}_t$ if $\Phi \subset \mathcal{M}_\mathcal{T}$. The valuation that is obtained in this way is an example of a *compound* valuation.

The dynamic valuations obtained from (6.1) are in general not strongly time consistent, even if all dynamic valuations in the collection Φ do have this property. Acceptance consistency holds, however; this follows from the lemma below.

LEMMA 6.1 *Let Φ be a collection of dynamic valuations $\phi = (\phi_t)_{t\in\mathcal{T}}$ in $\mathcal{M}_{\mathcal{T}}$. Define a dynamic valuation $\hat{\phi} = (\hat{\phi}_t)_{t\in\mathcal{T}}$ by (6.1). If all dynamic valuations $\phi \in \Phi$ are strongly time consistent, then the dynamic valuation $\hat{\phi}$ satisfies*

$$\hat{\phi}_s(X) \geq \hat{\phi}_s(\hat{\phi}_t(X))$$

for $X \in L^\infty$ and $s \leq t$.

PROOF We have $\hat{\phi}_s(X) = \operatorname{ess\,inf}_\phi \phi_s(X) = \operatorname{ess\,inf}_\phi \phi_s(\phi_t(X)) \geq \operatorname{ess\,inf}_\phi \phi_s(\operatorname{ess\,inf}_\phi \phi_t(X)) = \operatorname{ess\,inf}_\phi \phi_s(\hat{\phi}_t(X)) = \hat{\phi}_s(\hat{\phi}_t(X))$. □

It is possible to prove in fairly general situations that even conditional consistency holds. Let Φ_t denote the collection of $\mathcal{F}_t$-conditional evaluations obtained by taking the instantiations of the dynamic valuations in the collection Φ at time t:

$$\Phi_t := \{\psi \mid \psi = \phi_t \text{ for some } \phi \in \Phi\}. \tag{6.2}$$

We consider a setting in which the conditional expectation operator under the reference measure, E_s^P, acts as a 'central' element that can be combined with any element of Φ_t, i.e.,

$$\phi_t \in \Phi_t \Rightarrow E_s^P \phi_t \in \Phi_s \quad (s < t). \tag{6.3}$$

PROPOSITION 6.2 *Let $\Phi \subset \mathcal{M}_{\mathcal{T}}$ be a collection of strongly time consistent dynamic valuations that satisfies (6.3). Then the dynamic valuation $\hat{\phi}$ defined by (6.1) is conditionally consistent.*

PROOF Take s and t with $0 \leq s < t \leq T$, and let X be an element of L^∞. We need to show that $\hat{\phi}_t(X) \geq 0$ if and only if $\hat{\phi}_s(1_F X) \geq 0$ for all $F \in \mathcal{F}_t$. Suppose first that $\hat{\phi}_t(X) \geq 0$. For $F \in \mathcal{F}_t$, we then have, by Lemma 6.1, $\hat{\phi}_s(1_F X) \geq \hat{\phi}_s(\hat{\phi}_t(1_F X)) = \hat{\phi}_s(1_F \hat{\phi}_t(X)) \geq 0$.

Next, assume that $\hat{\phi}_t(X) \ngeq 0$. Then there must exist an element $\phi_t \in \Phi_t$, and $F \in \mathcal{F}_t$ such that $\phi_t(1_F X) \lneqq 0$, and, by assumption (6.3), $\hat{\phi}_s(1_F X) \leq E_s^P \phi_t(1_F X) \lneqq 0$. □

We prove sequential consistency under the following additional assumption, which describes a property of closure under conditional pasting:

$$\phi_t, \phi'_t \in \Phi_t \Rightarrow 1_F \phi_t + 1_{F^c} \phi'_t \in \Phi_t \quad (t \in \mathcal{T},\ F \in \mathcal{F}_t) \tag{6.4}$$

PROPOSITION 6.3 *Under the assumption (6.4) and the assumptions of Prop. 6.2, the dynamic valuation defined by (6.1) is sequentially consistent.*

PROOF We show that $\hat{\phi}_s(X) = 0$ for $s < t$ and $X \in L^\infty$ such that $\hat{\phi}_t(X) = 0$; sequential consistency then follows from Prop. 3.2. The inequality $\hat{\phi}_s(X) \geq 0$ follows from the assumption $\hat{\phi}_t(X) = 0$ by Lemma 6.1.

The reverse inequality can be proved as follows. From (6.4) it follows that for any pair ϕ_t, ϕ'_t in Φ_t, also their minimum $\min(\phi_t, \phi'_t)$ belongs to Φ_t, so that the family $\mathcal{E} := \{\phi_t(X) \mid \phi_t \in \Phi_t\}$ is directed downwards. It now follows (see for instance [14, Thm. A.33]) that there exists a sequence $(\phi_t^k)_{k \in \mathbb{N}}$ in Φ_t such that $\phi_t^k(X) \searrow 0$. Then, also for $\phi_s^k := E_s^P \phi_t^k$, $\phi_s^k(X) \searrow 0$. By assumption (6.3), $\phi_s^k \in \Phi_s$ for all k, and hence $\hat{\phi}_s(X) \leq 0$, which we had to prove. □

EXAMPLE 6.4 We consider compound valuations defined in terms of entropic risk measures on a discrete time axis $\mathcal{T} = \{0, \ldots, T\}$. Define, for a given $t \in \mathcal{T}$ and nonnegative, $\mathcal{F}_t$-measurable parameter β_t, the mapping $\bar{\phi}_t^{\beta_t} : L_{t+1}^\infty \to L_t^\infty$ by

$$\bar{\phi}_t^{\beta_t}(X) = -\frac{1}{\beta_t} \log E_t^P \exp(-\beta_t X) \qquad (X \in L_{t+1}^\infty). \tag{6.5}$$

For zero parameter values the right hand side is replaced by E_t^P, as usual. The mappings $\bar{\phi}_t^{\beta_t}$ can subsequently be pieced together to form monetary valuations:

$$\phi_T^{\beta_T}(X) = X \tag{6.6a}$$

$$\phi_t^{\beta_t}(X) = \bar{\phi}_t^{\beta_t}(\phi_{t+1}^{\beta_{t+1}}(X)). \tag{6.6b}$$

This recursion is a discrete-time version of the construction of families of conditional evaluations in terms of backward stochastic differential equations; cf. [6] and the references therein. The construction results in dynamic valuations $\phi^\beta \in \mathcal{M}_\mathcal{T}$, parametrized by the vector $\beta := (\beta_0, \beta_1, \ldots, \beta_{T-1})$. Note that by construction the valuations are strongly time consistent, unlike the ones obtained from a similar definition in [1, Section 4.5] in terms of adapted risk aversion levels over the entire remaining horizon. Compound valuations can now be defined by choosing sets $\mathcal{B}$ of admissible parameter vectors β, cf. (6.1). Strongly time consistent compound valuations are completely determined by their single-step properties, and hence correspond to letting $\mathcal{B}$ be equal to a singleton that specifies the applied level of risk aversion per time step. If one imposes that risk aversion is time- and state-independent, then more specifically $\beta_t = c$ for all $t \in \mathcal{T}$, where c is a fixed nonnegative real parameter. Sequentially consistent examples can be obtained, for instance, by considering sets of the form

$$\mathcal{B} = \{(\beta_0, \ldots, \beta_T) \mid \sum_{k=s}^{t-1} \beta_k \leq B(s,t)\}, \tag{6.7}$$

where B is a function that specifies a "budget" for risk tolerance across the interval $[s, t]$. This follows from the proposition above, which is applicable under the obvious condition that the budget does not decrease if the interval is extended. To illustrate that the additional freedom does not rely on introducing more parameters, one can consider taking $B(s,t) = c$. More generally, one can make the budget grow with period length, and set $B(s,t) = c(t-s)^\alpha$, for some $\alpha \in [0,1]$. This includes the strongly time consistent example given earlier, for $\alpha = 1$. By choosing lower values of α, one can tune long term features of the compound valuation *without changing the induced one-step properties*, a feature that is not available under strong time consistency.

We briefly indicate how this additional flexibility can be used to accommodate that two positions with the same ask price tomorrow may have different ask prices today, as argued in Example 3.9. Suppose that the bid and ask price of positions $X' \in L^\infty$ at t are given by resp. $b_t(X') = \hat{\phi}_t(X')$ and $a_t(X') = -\hat{\phi}_t(-X')$, where $\hat{\phi}$ is the dynamic compound valuation corresponding to a budget function of the form $B(s,t) = c(t-s)^\alpha$, as described above. For concreteness, consider a nonrecombining two-period binomial tree with uniform reference measure P, as in Example 4.5. Take $c = 1$, and consider the cases (i) $\alpha = 0$, (ii) $\alpha = \frac{1}{2}$, and (iii) the strongly time consistent case $\alpha = 1$. Following the idea of Example 3.9, consider position $X = (2,1,1,0)$, which yields $a_1(X) = (1.62, 0.62)$. So $Y = (1.62, 1.62, 0.62, 0.62)$ has the same ask price at $t = 1$, and its initial ask price turns out to be $a_0(Y) = 1.24$, independently of α. For X, however, the outcomes of $a_0(X)$ in the three cases are 1.12, 1.17, and, of course, 1.24, respectively. These values have been determined by first computing $-\bar{\phi}_1^{\beta_1}(-X)$ *as a (non-decreasing) function of* $\beta_1 \in [0,c]$, then applying $-\bar{\phi}_0^{\beta_0}$ to each outcome with β_0 equal to the maximum tolerated value $\min\{c, c2^\alpha - \beta_1\}$, and finally maximizing with respect to β_1. Note the additional degree of freedom in the backward recursion, related to keeping track of the amount of "consumed" risk aversion represented by β_1.

If one restricts attention to deterministic parameter values $\beta_t \in \mathbb{R}$, these constructions still yield conditionally consistent dynamic valuations, as claimed by Proposition 6.2. However, sequential consistency is lost in general if the number of time periods is more than two, as is easily verified by extending the example to three steps.

7 Conclusions

The construction of dynamic risk measures and nonlinear pricing operators that combine time consistency with reasonable levels of prudence across different time scales remains a challenging task. In this paper we have analyzed time consistency of families of monetary valuations, which are monetary risk measures under a positive sign convention. We have shown that, even under very weak interpretations of the notion of time consistency, the choice of the initial valuation already fully determines the family of valuations to which it belongs, if certain sensitivity conditions are satisfied. By considering different notions of time consistency, we obtain a categorization of monetary valuations in valuations that allow conditionally consistent updating, valuations that allow sequentially consistent updating, and valuations that allow strongly consistent updating. By means of examples related to regulation and to nonlinear pricing, we have indicated that weak forms of consistency can be motivated. We have provided characterizations of the possibility of updating in several senses, and we have given an example of the construction of consistent families which allow flexibility in the specification of prudence over time. An issue that calls for further research is the fact that the necessary and sufficient conditions that we have given for membership of consistent families are not always easy to verify. It would be desirable to have more readily verifiable necessary and/or sufficient conditions. Criteria for weak consistency in terms of dual representations are provided in [23].

8 Appendix

8.1 Auxiliary results

We prove a few auxiliary results. The following lemma can be proved by an argument similar to the reasoning in the proof of Thm. 4.33 in [14].

LEMMA 8.1 *Let $\phi : L^\infty \to L^\infty$ be a normalized monotonic mapping that is continuous from above, and let $X \in L^\infty$. If there exists a bounded sequence $(X_n)_{n\geq 1}$ such that $X_n \to X$ and $\phi(X_n) \geq 0$ for all n, then $\phi(X) \geq 0$.*

LEMMA 8.2 *Let $\phi : L^\infty \to L^\infty$ be strictly monotonic. If $X_t, Y_t \in L^\infty_t$ are such that $\phi(1_F X_t) \geq \phi(1_F Y_t)$ for all $F \in \mathcal{F}_t$, then $X_t \geq Y_t$.*

PROOF Let the assumptions of the lemma hold. If $X_t \ngeq Y_t$, then there exists $\varepsilon > 0$ such that the set $F = \{Y_t \geq X_t + \varepsilon\}$ has positive measure. It follows from $1_F X_t \leq 1_F(Y_t - \varepsilon)$ that

$$\phi(1_F Y_t) \leq \phi(1_F X_t) \leq \phi(1_F(Y_t - \varepsilon)) \leq \phi(1_F Y_t)$$

and consequently all inequalities above are in fact equalities. It then follows from the strong sensitivity of ϕ that $1_F = 0$, i. e. $P(F) = 0$, and we have a contradiction. □

COROLLARY 8.3 *Let $\phi : L^\infty \to L^\infty$ be strictly monotonic. If $X_t, Y_t \in L^\infty_t$ are such that $\phi(1_F X_t) = \phi(1_F Y_t)$ for all $F \in \mathcal{F}_t$, then $X_t = Y_t$.*

LEMMA 8.4 *Let $\phi : L^\infty \to L^\infty$ be normalized, monotonic, and sensitive. If $X_t \in L^\infty_t$ is such that $\phi(1_F X_t) \geq 0$ for all $F \in \mathcal{F}_t$, then $X_t \geq 0$.*

Proof Apply the argument in the proof of Lemma 8.2 with $Y_t = 0$, noting that $\phi(Y_t) = \phi(0) = 0$ and replacing strong sensitivity with sensitivity. $\square$

8.2 Proof of Prop. 2.2

The necessity of the five conditions has already been shown. To prove the remaining claims, we first have to show that (2.16) indeed defines a mapping $\phi^t_{\mathcal{S}} : L^\infty \to L^\infty_t$ if $\mathcal{S}$ satisfies the basic conditions. Take $X \in L^\infty$, and let $Y_t \in L^\infty_t$ be such that $X - Y_t \in \mathcal{S}$. It follows from the solidness of $\mathcal{S}$ that then also $\|X\|_t - Y_t \in \mathcal{S}$. Since $\|X\|_t - Y_t \in L^\infty_t$, the $\mathcal{F}_t$-nonnegativity of $\mathcal{S}$ implies that $Y_t \leq \|X\|_t$. This shows that the essential supremum in (2.16) is finite-valued (actually $\phi^t_{\mathcal{S}}(X) \leq \|X\|_t$) so that indeed $\phi^t_{\mathcal{S}}(X) \in L^\infty_t$ for every $X \in L^\infty$.

Next we verify that $\phi^t_{\mathcal{S}}$ has all the properties of an $\mathcal{F}_t$-conditional monetary valuation if $\mathcal{S}$ satisfies the basic conditions. The $\mathcal{F}_t$-nonnegativity of $\mathcal{S}$ and the assumption that $0 \in \mathcal{S}$ together imply that $\operatorname{ess\,inf} L^\infty_t \cap \mathcal{S} = 0$ so that $\phi^t_{\mathcal{S}}(0) = 0$ as required. The monotonicity property (2.2) of $\phi^t_{\mathcal{S}}$ is immediate from the solidness of $\mathcal{S}$. The conditional translation property (2.3) of $\phi^t_{\mathcal{S}}$ follows, in fact without any assumptions on the set $\mathcal{S}$, from the corresponding property of the essential supremum.

For the claim on $\mathcal{A}(\phi^t_{\mathcal{S}})$, we refer to [10, Prop. 3.10]. All claims but the last one now follow.

Finally we consider conditional concavity. It follows as in the proof of Prop. 3 in [12] that the convexity of $\mathcal{S}$ implies that the mapping $\phi^t_{\mathcal{S}}$ is $\mathcal{F}_0$-concave, i.e., $\phi^t_{\mathcal{S}}(\lambda X + (1-\lambda)Y \geq \lambda\phi^t_{\mathcal{S}}(X) + (1-\lambda)\phi^t_{\mathcal{S}}(Y)$ for $\lambda \in \mathbb{R}$ with $0 \leq \lambda \leq 1$ and $X, Y \in L^\infty$. It was shown in [10, Prop. 3.3] that monotonicity and conditional translation invariance of $\phi_t \in \mathcal{M}_t$ together imply the conditional local property as well as the inequality $\phi_t(X) - \phi_t(Y) \leq \|X - Y\|_t$ in L^∞_t, and in the same paper it is shown that the latter two properties together with $\mathcal{F}_0$-concavity imply $\mathcal{F}_t$-concavity [10, Rem. 3.4].

8.3 Proof of Prop. 3.6

The inclusion $\mathcal{A}_t \subset \mathcal{A}^t_s$ is equivalent to acceptance consistency as noted in the main text, and it is immediate from the definitions that acceptance consistency is implied by sequential consistency and by strong time consistency. It remains to prove that the reverse inclusion, consecutivity, holds under each of the three conditions mentioned. That is, we need to show the implication from right to left in (3.6). Take $X \in L^\infty$, and suppose that $\phi_s(1_F X) \geq 0$ for all $F \in \mathcal{F}_t$. Consider now each of the three conditions.

1. Take $F = \{\phi_t(X) \leq 0\}$, so that $1_F\phi_t(X) \leq 0$. Under strong time consistency, we can write

$$0 \leq \phi_s(1_F X) = \phi_s(\phi_t(1_F X)) = \phi_s(1_F \phi_t(X)) \leq 0.$$

It follows that $\phi_s(1_F\phi_t(X)) = 0$; sensitivity then implies that $1_F\phi_t(X) = 0$, or in other words $\phi_t(X) \geq 0$.

2. Take F as above. Under sequential consistency we can write, using item (ii) in Lemma 3.2:

$$0 \leq \phi_s(1_F X) \leq \phi_s(1_F X - 1_F\phi_t(X)) = \phi_s(1_F X - \phi_t(1_F X)) = 0. \tag{8.1}$$

It follows that all inequalities are in fact equalities, and strong sensitivity of ϕ_s implies that $1_F\phi_t(X) = 0$, i.e. $\phi_t(X) \geq 0$.

3. Strict sequential consistency implies that, for $X \in L^\infty$ such that $\phi_t(X) \leq 0$, we can conclude that $\phi_t(X) = 0$ when $\phi_s(X) = 0$. With F as above, the relations (8.1) imply that $\phi_s(1_F X) = 0$ whereas we also have $\phi_t(1_F X) = 1_F\phi_t(X) \leq 0$. It follows once more that $1_F\phi_t(X) = 0$.

8.4 Proof of Prop. 3.7

Concerning item (i), since ϕ_s is sensitive, its acceptance set has the negative cone exclusion property. This property is inherited by the $\mathcal{F}_t$-restriction of $\mathcal{A}(\phi_s)$ which is the acceptance set of ϕ_t, and it follows that ϕ_t is sensitive as well. Item (ii) follows from Prop. 2.2.

Now consider the claim concerning continuity. Suppose that $(X_n)_{n\geq 1}$ is a nonincreasing sequence of elements of L^∞ that converges to $X \in L^\infty$. By the monotonicity of ϕ_t, the sequence $(\phi_t(X_n))_{n\geq 1}$ is nonincreasing as well and is bounded from below by $\phi_t(X)$, so that we can define $Z = \lim_{n\to\infty} \phi_t(X_n)$. To prove the continuity from above, we must show that $Z = \phi_t(X)$. We have $\phi_t(X_n) \geq \phi_t(X)$ for all n, which already implies that $Z \geq \phi_t(X)$. Because Z is the pointwise limit of a sequence of $\mathcal{F}_t$-measurable functions, it is itself $\mathcal{F}_t$-measurable, so the inequality $Z \leq \phi_t(X_n)$, which holds for each n, may be written as $\phi_t(X_n - Z) \geq 0$. By conditional consistency, this means that $\phi_s(1_F(X_n - Z)) \geq 0$ for all $F \in \mathcal{F}_t$. Since $1_F(X_n - Z) \searrow 1_F(X - Z)$, the assumed continuity from above of ϕ_s implies that $\phi_s(1_F(X - Z)) \geq 0$ for all $F \in \mathcal{F}_t$, which means that $\phi_t(X - Z) \geq 0$. Again using the $\mathcal{F}_t$-measurability of Z, we conclude that $\phi_t(X) \geq Z$.

8.5 Proof of Prop. 4.6

The property expressed in Prop. 4.6 for sequentially or strongly time consistent updating does not follow from the corresponding property for conditionally consistent updating, since we need to prove an implication that has a stronger premise but also a stronger conclusion. We therefore provide three separate proofs.

Conditional consistency

Assume that ϕ_u is a conditionally consistent update of ϕ_t and ϕ_t is a conditionally consistent update of ϕ_s. We want to show that ϕ_u is also a conditionally consistent update of ϕ_s, which means that $\phi_u(X) \geq 0$ if and only if $\phi_s(1_F X) \geq 0$ for all $F \in \mathcal{F}_u$. First, take $X \in L^\infty$ such that $\phi_u(X) \geq 0$. For all $F \in \mathcal{F}_u$, we have $\phi_t(1_F X) \geq 0$ which implies that $\phi_s(1_F X) \geq 0$. Conversely, suppose that $\phi_s(1_F X) \geq 0$ for all $F \in \mathcal{F}_u$. Take $F \in \mathcal{F}_u$ and $F' \in \mathcal{F}_t \subset \mathcal{F}_u$; then, since $F' \cap F \in \mathcal{F}_u$, we have $\phi_s(1_{F'} 1_F X) = \phi_s(1_{F' \cap F} X) \geq 0$. The fact that this holds for all $F' \in \mathcal{F}_t$ implies, because ϕ_t is a conditionally consistent update of ϕ_s, that $\phi_t(1_F X) \geq 0$. This inequality in its turn holds for all $F \in \mathcal{F}_u$, and so, because ϕ_u is a conditionally consistent update of ϕ_t, it follows that $\phi_u(X) \geq 0$.

Now assume that both ϕ_u and ϕ_t are conditionally consistent updates of ϕ_s. We want to show that $\phi_u(X) \geq 0$ if and only if $\phi_t(1_F X) \geq 0$ for all $F \in \mathcal{F}_u$. First, take $X \in L^\infty$ such that $\phi_u(X) \geq 0$. Take $F \in \mathcal{F}_u$. For all $F' \in \mathcal{F}_t$ we have $F \cap F' \in \mathcal{F}_u$ so that $\phi_s(1_{F'} 1_F X) \geq 0$. It follows that $\phi_t(1_F X) \geq 0$. Conversely, suppose that $\phi_t(1_F X) \geq 0$ for all $F \in \mathcal{F}_u$; then we also have $\phi_s(1_F X) \geq 0$ for all $F \in \mathcal{F}_u$, so that $\phi_u(X) \geq 0$.

(Strict) sequential consistency

Under strong sensitivity, the strict and non-strict versions are equivalent, cf. (3.2). We formulate the proof for sequential consistency. If ϕ_u is a sequentially consistent update of ϕ_t, then it follows immediately from the definition that it is also a sequentially consistent update of ϕ_s. Assume now that ϕ_u is a sequentially consistent update of ϕ_s, and suppose it is not a sequentially consistent update of ϕ_t, due to a violation of acceptance consistency (3.1a) (the proof in case of a rejection inconsistency is analogous). Then there exists $X \in L^\infty$ such that $\phi_u(X) \geq 0$ and $\phi_t(X) \not\geq 0$, so that there is an $F \in \mathcal{F}_t$ with $P(F) > 0$ and an $\varepsilon > 0$ such that $1_F \phi_t(X) \leq -\varepsilon 1_F$. Take $\eta \in (0, \varepsilon)$. Because $F \in \mathcal{F}_t \subset \mathcal{F}_u$, we have $\phi_u(1_F(X + \eta)) \geq \phi_u(1_F(X)) = 1_F \phi_u(X) \geq 0$ so that $\phi_s(1_F(X + \eta)) \geq 0$ by the assumed sequential consistency of ϕ_u and ϕ_s. The conditional monetary valuation ϕ_t is also a sequentially consistent update of ϕ_s, so that from $\phi_t(1_F(X + \eta)) = 1_F(\phi_t(X) + \eta) \leq 0$ it follows that $\phi_s(1_F(X + \eta)) \leq 0$. We conclude that $\phi_s(1_F(X + \eta)) = 0$. Since this holds for all $0 < \eta < \varepsilon$, strong sensitivity of ϕ_s now implies that $1_F = 0$, and we have a contradiction.

Strong time consistency

Assume that ϕ_u is a strongly time consistent update of ϕ_t, and ϕ_t of ϕ_s. Then, for any $X \in L^\infty$, we have $\phi_s(\phi_u(X)) = \phi_s(\phi_t(\phi_u(X))) = \phi_s(\phi_t(X)) = \phi_s(X)$ so that ϕ_u is a strongly time consistent update of ϕ_s. Conversely, assume now that ϕ_t and ϕ_u are strongly time consistent updates of ϕ_s. Take $X \in L^\infty$. For any $F \in \mathcal{F}_t$, we have, since $\mathcal{F}_t \subset \mathcal{F}_u$,

$$\phi_s(1_F\phi_t(\phi_u(X))) = \phi_s(\phi_t(\phi_u(1_F X))) = \phi_s(\phi_u(1_F X)) = \phi_s(1_F X) = \phi_s(1_F\phi_t(X))$$

and it follows that $\phi_t(\phi_u(X)) = \phi_t(X)$ by Cor. 8.3 in the Appendix. This corollary applies if ϕ_t is strongly sensitive. This follows from the following lemma, which is applicable because sequential consistency is implied by strong time consistency.

LEMMA 8.5 *A sequentially consistent update of a strongly sensitive conditional monetary valuation is itself strongly sensitive.*

PROOF Let $\mathcal{F}_s$ and $\mathcal{F}_t$ be sub-σ-algebras such that $\mathcal{F}_s \subset \mathcal{F}_t$; let $\phi_s \in \mathcal{M}_s$ and $\phi_t \in \mathcal{M}_t$ Suppose that ϕ_s is strongly sensitive and that ϕ_t is a sequentially consistent update of ϕ_s. To prove that ϕ_t is strongly sensitive as well, take $X, Y \in L^\infty$ such that $X \geq Y$ and $\phi_t(X) = \phi_t(Y)$. Due to sequential consistency (cf. item (ii) in Lemma 3.2), we have $\phi_s(X - \phi_t(X)) = 0$ and also $\phi_s(Y - \phi_t(X)) = \phi_s(Y - \phi_t(Y)) = 0$. Since $X - \phi_t(X) \geq Y - \phi_t(X)$, strong sensitivity of ϕ_s implies that $X - \phi_t(X) = Y - \phi_t(X)$ and therefore $X = Y$. □

8.6 Proof of Lemma 5.2

Write $\mathcal{S} := \mathcal{A}(\phi)$. Take $X \in L^\infty$, and let $(X_n)_{n\geq 1}$ be a sequence of payoffs $X_n \in \mathcal{S}^t$ such that $\|X_n - X\|_t \to 0$. Note that we then also have $X_n \to X$. Take $F \in \mathcal{F}_t$; we want to show that $1_F X \in \mathcal{S}$. By Egorov's theorem, we can find for any $m \in \mathbb{N}$ a set $G_m \in \mathcal{F}_t$ with $P(G_m) > 1 - \frac{1}{m}$ such that the convergence of $\|X_n - X\|_t$ to 0 is uniform on G_m. In particular it follows, for fixed m, that $(1_{G_m}\|X_n - X\|_t)_{n\geq 1}$ is a bounded sequence, which implies that $(1_{G_m} X_n)_{n\geq 1}$ is a bounded sequence as well. From $X_n \to X$ it follows that $1_{G_m\cap F} X_n \to 1_{G_m \cap F} X$. Moreover, since $G_m \in \mathcal{F}_t$ and $X_n \in (\mathcal{A}(\phi))^t$, we have $\phi(1_{G_m\cap F} X_n) \geq 0$ for all n. By Lemma 8.1, it follows that $\phi(1_{G_m \cap F} X) \geq 0$. Now, the sequence $(1_{G_m\cap F} X)_{m\geq 1}$ is a bounded sequence that converges to $1_F X$ and that satisfies $\phi(1_{G_m\cap F} X) \geq 0$ for all m. Using Lemma 8.1 again, we conclude that $\phi(1_F X) \geq 0$. Since $F \in \mathcal{F}_t$ was arbitrary, it follows that $X \in \mathcal{S}^t$.

8.7 Proof of Prop. 5.6

If ϕ_s has a sequentially consistent update ϕ_t, then $C_t = \phi_t(X)$ satisfies the requirements of the proposition. Conversely, suppose now that $\phi_s \in \mathcal{M}_s$ has a conditionally consistent $\mathcal{F}_t$-update, say ϕ_t, and that for each $X \in L^\infty$ there exists $C_t \in L^\infty_t$ such that (5.6) holds. To prove that the update ϕ_t is sequentially consistent, it is sufficient, in view of Lemma 3.2, to show that the latter condition implies $C_t = \phi_t(X)$. Therefore, take $X \in L^\infty$, and let

$C_t \in L_t^\infty$ be such that (5.6) holds. By conditional consistency, the condition (5.6) implies that $\phi_t(X - C_t) \geq 0$ and hence $C_t \leq \phi_t(X)$. To prove the reverse inequality, take $Y_t \in L_t^\infty$ and suppose that

$$\phi_s(1_F(X - C_t - Y_t)) \geq 0 \quad \text{for all } F \in \mathcal{F}_t.$$

Take in particular $F = \{Y_t \geq 0\}$. We then have $1_F Y_t \geq 0$ so that $1_F(X - C_t) \geq 1_F(X - C_t - Y_t)$. Using (5.6), we can write

$$0 = \phi_s(1_F(X - C_t)) \geq \phi_s(1_F(X - C_t - Y_t)) \geq 0.$$

The strong sensitivity of ϕ_s now implies that $1_F Y_t = 1_{Y_t \geq 0} Y_t = 0$ so that $Y_t \leq 0$. We have shown that

$$\operatorname{ess\,sup}\{Y_t \in L_t^\infty \mid \phi_s(1_F(X - C_t - Y_t)) \geq 0 \text{ for all } F \in \mathcal{F}_t\} \leq 0. \tag{8.2}$$

The conditional monetary valuation ϕ_t must be equal to the refinement update of ϕ_s, by Thm. 4.4. In view of the expression given for the refinement update in (4.1), it follows from (8.2) that $\phi_t(X - C_t) \leq 0$. Therefore, we obtain the inequality $\phi_t(X) \leq C_t$, and the proof is complete.

8.8 Proof of Prop. 5.7

If ϕ_s admits a strongly time consistent update ϕ_t, then $C_t = \phi_t(X)$ satisfies the requirements of the proposition; indeed, $\phi_t(X) \in L_t^\infty$ and, for all $F \in \mathcal{F}_t$, $\phi_s(1_F \phi_t(X)) = \phi_s(\phi_t(1_F X)) = \phi_s(1_F X)$. Conversely, suppose now that for each $X \in L^\infty$ there exists $C_t \in L_t^\infty$ such that (5.7) holds. It follows from Cor. 8.3 that for each given X there can be only one such $C_t \in L_t^\infty$, and so we can define a mapping $\psi : L^\infty \to L_t^\infty$ implicitly by

$$\phi_s(1_F X) = \phi_s(1_F \psi(X)) \quad (F \in \mathcal{F}_t). \tag{8.3}$$

If we can show that the mapping ψ belongs to $\mathcal{M}_t$, then strong time consistency follows from (8.3) and the proof will be complete.

In order to prove that $\psi \in \mathcal{M}_t$, it suffices [9, Rem. 3.4] to prove that ψ is normalized and monotonic, and that it satisfies the local property as well as translation invariance (i.e. $\psi(X + m) = \psi(X) + m$ for $X \in L^\infty$ and $m \in \mathbb{R}$). The normalization property is trivial, and monotonicity follows from an application of Lemma 8.2 in the Appendix. Because ψ is normalized, the local property is equivalent to regularity. Take $G \in \mathcal{F}_t$ and $X \in L^\infty$. We have, for all $F \in \mathcal{F}_t$,

$$\phi_s(1_F \psi(1_G X)) = \phi_s(1_F 1_G X) = \phi_s(1_F 1_G \psi(X))$$

and moreover $1_G \psi(X) \in L_t^\infty$, so that $\psi(1_G X) = 1_G \psi(X)$ as was to be proved. To show translation invariance, first note that $\psi(m) = m$ for all $m \in \mathbb{R}$. Now take $X \in L^\infty$ and

$m \in \mathbb{R}$. Using the translation invariance of ϕ_s as well as the regularity property of ψ which has already been proved and the property $\phi_s(X) = \phi_s(\psi(X))$ which is a special case of (8.3), we can write, for $F \in \mathcal{F}_t$,

$$\begin{aligned}\phi_s(1_F(X+m)) &= \phi_s(1_F X - 1_{F^c} m) + m = \phi_s(\psi(1_F X - 1_{F^c} m)) + m = \\ &= \phi_s(1_F \psi(X) - 1_{F^c} m) + m = \phi_s(1_F \psi(X) + 1_F m) = \\ &= \phi_s(1_F(\psi(X) + m)).\end{aligned}$$

Also, we have $\psi(X) + m \in L_t^\infty$. It follows that $\psi(X+m) = \psi(X) + m$, and this completes the proof.

References

[1] B. Acciaio and I. Penner. Dynamic risk measures. In G. Di Nunno and B. Øksendal, editors, *Advanced Mathematical Methods for Finance*, pages 1–34. Springer, Heidelberg, 2011.

[2] Ph. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9:203–228, 1999.

[3] Ph. Artzner, F. Delbaen, J.-M. Eber, D. Heath, and H. Ku. Coherent multiperiod risk adjusted values and Bellman's principle. *Annals of Operations Research*, 152:5–22, 2007.

[4] P. Barrieu and N. El Karoui. Inf-convolution of risk measures and optimal risk transfer. *Finance and Stochastics*, 9:269–298, 2005.

[5] V. Bignozzi and A. Tsanakas. Characterization and construction of sequentially consistent risk measures. Working paper, Cass Business School, City University London, 2012.

[6] J. Bion-Nadal. Dynamic risk measures: time consistency and risk measures from BMO martingales. *Finance and Stochastics*, 12:219–244, 2008.

[7] C. Burgert. *Darstellungssätze für statische und dynamische Risikomaße mit Anwendungen.* PhD thesis, Universität Freiburg, 2005.

[8] P. Carr, H. Geman, and D.B. Madan. Pricing and hedging in incomplete markets. *Journal of Financial Economics*, 32:131–167, 2001.

[9] P. Cheridito, F. Delbaen, and M. Kupper. Coherent and convex monetary risk measures for unbounded càdlàg processes. *Finance and Stochastics*, 9:369–387, 2005.

[10] P. Cheridito, F. Delbaen, and M. Kupper. Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability*, 11:57–106, 2006.

[11] A. Cherny and D.B. Madan. Markets as a counterparty: an introduction to conic finance. *International Journal of Theoretical and Applied Finance*, 13:11491177, 2010.

[12] K. Detlefsen and G. Scandolo. Conditional and dynamic convex risk measures. *Finance and Stochastics*, 9:539–561, 2005.

[13] H. Föllmer and I. Penner. Convex risk measures and the dynamics of their penalty functions. *Statistics and Decisions*, 24:61–96, 2006.

[14] H. Föllmer and A. Schied. *Stochastic Finance. An Introduction in Discrete Time* (3rd ed.). Walter de Gruyter, Berlin, 2011.

[15] M. Frittelli and E. Rosazza Gianin. Putting order in risk measures. *Journal of Banking and Finance*, 26:1473–1486, 2002.

[16] M. Frittelli and E. Rosazza Gianin. Dynamic convex risk measures. In G. Szegö, editor, *Risk Measures for the 21st Century*, pages 227–248. Wiley, New York, 2004.

[17] A. Jobert and L.C.G. Rogers. Valuations and dynamic convex risk measures. *Mathematical Finance*, 18:1–22, 2008.

[18] S. Klöppel and M. Schweizer. Dynamic utility indifference valuation via convex risk measures. NCCR FINRISK Working Paper 209, ETH Zürich, 2005.

[19] M. Kupper and W. Schachermayer. Representation results for law invariant time consistent functions. *Mathematics and Financial Economics*, 3:189–210, 2009.

[20] I. Penner. *Dynamic Convex Risk Measure: Time Consistency, Prudence, and Sustainability.* PhD thesis, Humboldt-Universität, Berlin, 2007.

[21] F. Riedel. Dynamic coherent risk measures. *Stochastic Processes and their Applications*, 112:185–200, 2004.

[22] 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.

[23] B. Roorda and J.M. Schumacher. Weakly consistent convex risk measures and their dual representations. Working Paper, 2012.

[24] B. Roorda, J.M. Schumacher, and J.C. Engwerda. Coherent acceptability measures in multiperiod models. *Mathematical Finance*, 15:589–612, 2005.

[25] A. Schied. Optimal investments for risk- and ambiguity-averse preferences: a duality approach. *Finance and Stochastics*, 11:107–129, 2007.

[26] S. Tutsch. *Konsistente und konsequente dynamische Risikomaße und das Problem der Aktualisierung*. PhD thesis, Humboldt-Universität, Berlin, 2006.

[27] S. Tutsch. Update rules for convex risk measures. *Quantitative Finance*, 8:833–843, 2008.

[28] S. Weber. Distribution-invariant risk measures, information, and dynamic consistency. *Mathematical Finance*, 16:419–441, 2006.