Membership Conditions for Consistent Families of Risk Measures

Berend Roorda*   Hans Schumacher†

May 7, 2012

Abstract

Through a number of recent research contributions, it has become clear that time consistency imposes strong constraints on families of risk measures that are designed to operate on different time horizons. In this paper we add to these results by showing that consistent families of risk measures are already determined uniquely by the choice of the risk measure at the highest level of aggregation; moreover, this statement holds even when the term "consistency" is interpreted in a rather weak sense. The unique update rule is specified explicitly. We then derive conditions that must be satisfied for risk measures to belong to consistent families. An application is given to the construction of consistent families of compound risk measures.

Keywords: risk measures; acceptability functionals; updating; weak time consistency.

1 Introduction

Risk measures are used for various purposes, including regulation, margin setting, asset pricing, and contract design; see for instance [2, 8, 4, 12]. In many applications, it is of interest to carry over risk measures from earlier to later times or more generally from higher to lower levels of aggregation. Updating rules and time consistency of convex and coherent risk measures have been investigated extensively in recent years; see for instance [16, 28, 31, 13, 10, 3]. Attention has been paid in particular to the notion of *dynamic* or

*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.

strong consistency, which is standard in linear pricing theory. Weaker notions of consistency are supported by the interpretation in which the amount computed from a risk measure is taken as a *reserve* that should be maintained in relation to a certain position, rather than as a price. As we have argued before [29], the recursive nature of strong consistency makes it hard to develop families of risk measures that maintain comparable standards of prudence on different horizons; a "VaR of VaR's", so to say, is likely to be very conservative. For related comments see also Schied [32, Rem. 3.5].

The fact that dynamic consistency is a strong requirement has been noted in the literature on statistical decision theory. Machina [25], and Hanany and Klibanoff [18] show that it is problematic to combine various forms of dynamic consistency with preferences that are not of the expected utility type. The restrictive nature of dynamic consistency is reflected in the fact that many risk measures do not allow updates that are consistent in this sense. Klöppel and Schweizer [23, Section 7.2] introduce a 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 [24] shows that, under mild technical conditions, law-invariant risk measures allow dynamically 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, 29, 35, 34, 1, 5]. The main notions used in this paper are *sequential consistency* and *conditional consistency*. Both consistency notions have been introduced in [29] 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 intuitive 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 form of consistency, which is weaker than sequential consistency, updates are unique.

The uniqueness of updates 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 use of risk measures to define bid and ask prices has generated recent interest; we discuss the relations between time consistency and absence of arbitrage in this context. The paper concludes with an example of the construction of consistent families.

In this paper we consider the evaluation of *payoffs* (random variables) rather than of *payoff streams* (random processes) as for instance in [10] and [19]. 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 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 [19]. Their alternative term “valuation” perhaps points too strongly in the direction of pricing. To have a term that may cover a market price as well as an amount of regulatory capital, we add one letter and use the term “evaluation” instead, as has been done before by Peng [26]. 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 evaluation functionals correspond to acceptable positions.

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 is followed by a section in which conditions for absence of arbitrage are discussed in the setting in which risk measures are used to define bid and ask prices. The refinement update is introduced in Section 5 as the unique candidate for a consistent update. Existence of such an update is addressed in Section 6. An example of the construction of a consistent family is shown in Section 7. Finally, conclusions follow in Section 8. 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 [13, 10, 15]).

2.1 Standing assumptions and notation

We work in the standard setting of a 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 $P(X \leq Y) = 1$ 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 [15]); $\operatorname{ess\,inf} \mathcal{S}$ 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 $E_t^Q X$ rather than as $E^Q[X \mid \mathcal{F}_t]$.

Given a random variable $X \in L^\infty$, the variable $\|X\|_t \in L_t^\infty$ defined by $\|X\|_t = \operatorname{ess\,inf}\{m \in L_t^\infty \mid m \geq |X|\}$ 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 $L_t^\infty \subset L^\infty$ and $\|X\|_t \leq \|X\|$ for all $X \in L^\infty$.

A subset $\mathcal{S}$ of L^∞ will be called *real-convex* if $\lambda X + (1-\lambda)Y \in \mathcal{S}$ for all $X, Y \in \mathcal{S}$ and $\lambda \in \mathbb{R}$ such that $0 \leq \lambda \leq 1$. This is of course the usual notion of convexity, but we want to have a term that emphasizes the difference with the notion of $\mathcal{F}_t$-convexity defined in (2.18) below.

While the usual interpretation of a filtration member $\mathcal{F}_t$ is that of representing information available at time t, there is nothing in the developments below that prevents other possible interpretations, for instance information available to a particular agent, as mentioned in [13]. The σ-algebra $\mathcal{F}_t$ may also represent information which is available in principle but which may be used or not used at the discretion of a regulatory authority, or it may represent a certain aggregation level within an organization.

2.2 Conditional evaluations

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

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

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

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

The term "unconditional evaluation" is sometimes used for a mapping that satisfies the above properties with $t = 0$, so that effectively the mapping is from L^∞ to $\mathbb{R}$. An $\mathcal{F}_t$-conditional evaluation is said to be *concave* if it satisfies $\mathcal{F}_t$*-concavity*:

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

A concave $\mathcal{F}_t$-conditional evaluation is called *coherent* if it satisfies $\mathcal{F}_t$*-positive homogeneity*:

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

An $\mathcal{F}_t$-conditional evaluation always satisfies the $\mathcal{F}_t$*-local property* [13, 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* [13, 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):

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

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

A mapping that is monotonic and strongly sensitive is said to be *strictly monotonic*. A conditional evaluation 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 evaluation* corresponding to the filtration $(\mathcal{F}_t)_{t\in\mathcal{T}}$ is a family $(\phi_t)_{t\in\mathcal{T}}$ of mappings such that, for each $t \in \mathcal{T}$, ϕ_t is an $\mathcal{F}_t$-conditional evaluation.

2.3 Acceptance sets and conditional requirements

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

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

The acceptance set of an $\mathcal{F}_t$-conditional evaluation 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". The acceptance set of an $\mathcal{F}_t$-conditional evaluation always has the $\mathcal{F}_t$*-local property* (2.14) and the $\mathcal{F}_t$*-closedness* property (2.15):

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

[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 $\mathcal{F}_t$.

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

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 evaluation. The proposition below, obtained from [13] and [10], states this fact and also explains how to relate an $\mathcal{F}_t$-conditional evaluation to a subset $\mathcal{S}$ satisfying only the basic conditions. That construction relies on the notion of a *conditional capital requirement*, introduced in [13], which associates to an arbitrary set $\mathcal{S} \subset L^\infty$ the mapping from L^∞ to L^∞_t given by

$$\phi^t_{\mathcal{S}}(X) = \operatorname{ess\,sup}\{Y_t \in L^\infty_t \mid X - Y_t \in \mathcal{S}\}. \tag{2.16}$$

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

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

$$\phi_t = \phi^t_{\mathcal{A}(\phi_t)}. \tag{2.17}$$

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

$$\Lambda_t X + (1-\Lambda_t)Y \in \mathcal{S} \quad (X, Y \in \mathcal{S};\ \Lambda_t \in L^\infty_t,\ 0 \leq \Lambda_t \leq 1). \tag{2.18}$$

Sensitivity of a conditional evaluation 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 ϕ_s and ϕ_t be conditional evaluations with respect to $\mathcal{F}_s$ and $\mathcal{F}_t$, respectively, 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, 14, 33] while the property (3.1b) has been called "weak rejection consistency" [33]. 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 [35].

The following characterizations of sequential consistency may aid the intuition (cf. [29, Thm. 4.2], [33, 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 conditional evaluation ϕ_t is a sequentially consistent update of ϕ_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

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

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

3.2 Strong 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], [17, Def. 18], [14, Def. 3.1]. We will also refer to it as *strong consistency*.

DEFINITION 3.3 Let ϕ_s and ϕ_t be conditional evaluations with respect to $\mathcal{F}_s$ and $\mathcal{F}_t$, respectively, with $s \le t$. We say that ϕ_s and ϕ_t are *strongly time consistent*, or that ϕ_t is a *strongly consistent update* of ϕ_s, if the following relation holds for all $X \in L^\infty$:

$$\phi_s(\phi_t(X)) = \phi_s(X). \tag{3.2}$$

A characterization in terms of acceptance sets is given in [14]. The definition of strong consistency is sometimes given in the form of an implication: $\phi_t(X) = \phi_t(Y) \Rightarrow \phi_s(X) = \phi_s(Y)$ for $X, Y \in L^\infty$. Under $\mathcal{F}_t$-translation invariance, this is equivalent to the definition above, as can be seen by taking $Y = \phi_t(X)$. It is immediately clear from the definitions that strong consistency implies sequential consistency.

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 ϕ_s and ϕ_t be conditional evaluations with respect to $\mathcal{F}_s$ and $\mathcal{F}_t$, respectively, 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:

$$\phi_t(X) \ge 0 \;\Leftrightarrow\; \forall F \in \mathcal{F}_t : \phi_s(1_F X) \ge 0 \quad (X \in L^\infty). \tag{3.3}$$

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 *$\mathcal{F}_t$-refinement* of $\mathcal{S}$ is the set $\mathcal{S}^t$ defined by

$$\mathcal{S}^t = \{X \in \mathcal{S} \,|\, 1_F X \in \mathcal{S} \text{ for all } F \in \mathcal{F}_t\}. \tag{3.4}$$

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 [33, p. 88], in the situation in which the set $\mathcal{S}$ is the acceptance set $\mathcal{A}_s$ of a conditional evaluation ϕ_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

$$\mathcal{A}_t = \mathcal{A}_s^t. \tag{3.5}$$

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 (3.5). The reverse inclusion 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. The notion of *strict sequential consistency* used below is defined as sequential consistency with the added requirement $\phi_t(X) \lneqq 0 \Rightarrow \phi_s(X) \lneqq 0$ for $X \in L^\infty$.

Proposition 3.6 *Let ϕ_s and ϕ_t be conditional evaluations with respect to $\mathcal{F}_s$ and $\mathcal{F}_t$ respectively, 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 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 ϕ_s be a sensitive $\mathcal{F}_s$-conditional evaluation 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 9.5) that strong sensitivity is preserved under sequentially consistent updating.

4 Consistency and absence of arbitrage

The usual interpretation of risk measures is that they express a reserve capital that should be maintained in connection with a given risky position, in order to ensure solvency even under unfavorable conditions. Such an interpretation is associated with a strong or even exclusive focus on the tail of the loss distribution. In the axioms that are typically used, the focus on losses is reflected in the convexity axiom, but only in a weak sense. Indeed, linear expectation operators do satisfy the usual axioms, but they need not imply a special emphasis on negative outcomes as opposed to positive outcomes. The axioms may therefore be associated to more interpretations than the one that is suggested by the term "risk measure", which is one of the reasons why in this paper we use the more neutral term "evaluation" instead. In particular, evaluation functionals can be used to model bid and ask prices. The question may then be asked how consistency of dynamic evaluations relates to absence of arbitrage, and in particular whether strong consistency is necessary for absence of arbitrage. We provide an answer to that question in this section.

Consider an economy in which all uncertainty will be resolved at time T. The assets in the economy are contracts for delivery of a bounded contingent cashflow at time T. The interest rate is assumed to be zero, so that cashflows that are resolved at times $t < T$ are equivalent to cashflows at time T. Suppose that, in this economy, bid and ask prices are given by a dynamic evaluation $\phi = (\phi_t)_{t \in \mathcal{T}}$. Specifically, the ask price $a_t(X)$ and the bid price $b_t(X)$ at time $0 \le t \le T$ of the asset that delivers the cashflow X at time T are given by

$$a_t(X) = -\phi_t(-X), \qquad b_t(X) = \phi_t(X). \tag{4.1}$$

The bid-ask spread for an asset X is zero in case $\phi_t(-X) = -\phi_t(X)$ for all $t \in \mathcal{T}$. If this holds for all assets, then we say that the Law of One Price holds. The use of bid-ask spreads of this form goes back at least to Jouini and Kallal [20], who worked with sublinear functionals rather than with risk measures; that is, they used only the axioms of homogeneity and subadditivity.

In this section we will in particular consider coherent conditional evaluations. Under the coherence axioms, the set $\mathcal{A}_t$ given by

$$\mathcal{A}_t := \{X \in L^\infty \mid \phi_t(X) \ge 0\}$$

is a convex cone. This set is known under various names in the literature, such as *acceptance set* [2], *solvency region* [21], *acceptable opportunity set* [8], and *cone of marketed cash flows* [12]. We shall refer to it as the acceptance cone. In the pricing interpretation, the acceptance cone can be thought of as consisting of the positions that can be liquidated at no cost at time t, or in other words positions that the market at time t is willing to accept. A portfolio trading strategy will be said to be *budget feasible*, or simply *feasible*, if at any rebalancing time the change of composition of the portfolio is such that the market is willing to take the opposite position. In other words, the portfolio composition at time T that results from a feasible strategy is given by

$$X_T = X_0 + \sum_{j=0}^{N} \Delta X_{t_j}, \qquad -\Delta X_{t_j} \in \mathcal{A}_{t_j} \quad (j = 1, \ldots, N) \tag{4.2}$$

where $0 \le t_0 \le \cdots \le t_N \le T$, and conversely any position of this form can be reached by a feasible strategy starting from the initial portfolio X_0. The rule "$dX_t \in -\mathcal{A}_t$" was proposed by Kabanov [21]. Here we skirt the issue of finding a measurable implementation of the trading strategy. The market with bid-ask prices determined by the dynamic evaluation ϕ is said to *admit arbitrage* if there exists a position X_T of the form (4.2), with $X_0 = 0$, such that $X_T \gneqq 0$. Various alternative formulations are possible; cf. for instance [22].

Lemma 4.1 *A necessary condition for absence of arbitrage is that, for all $0 \le s \le t \le T$, the following inequality holds for all $X \in L^\infty$:*

$$\phi_s(X) \le -\phi_s(\phi_t(-X)). \tag{4.3}$$

Proof Let time points s and t be given, and take $X \in L^\infty$. For any $F \in \mathcal{F}_s$, we can define

$$\Delta X_s = 1_F\big(\phi_s(X) - X + \phi_s(\phi_t(-X)) - \phi_t(-X)\big)$$

$$\Delta X_t = 1_F\big(\phi_t(-X) + X\big).$$

Given any position X and $F \in \mathcal{F}_s$, we have $\phi_s\big(1_F(X - \phi_s(X))\big) \geq 0$, so that $1_F\big(\phi_s(X) - X\big) \in -\mathcal{A}_s$. Application of this rule both to X and to $\phi_t(-X)$ leads, together with the superadditivity of ϕ_s, to the conclusion that $\Delta X_s \in -\mathcal{A}_s$. Likewise, it follows that $\Delta X_t \in -\mathcal{A}_t$. Consequently, the position

$$\Delta X_s + \Delta X_t = 1_F\left(\phi_s(X) + \phi_s(\phi_t(-X))\right)$$

is reachable by a feasible strategy from the zero initial position. Now define F by

$$F = \{X \mid \phi_s(X) + \phi_s(\phi_t(-X)) > 0\}$$

and note that indeed $F \in \mathcal{F}_s$. To prevent arbitrage, we must have $1_F\big(\phi_s(X)+\phi_s(\phi_t(-X))\big) = 0$ or in other words

$$\phi_s(X) + \phi_s(\phi_t(-X)) \leq 0.$$

This is what we needed to prove. □

In terms of bid and ask prices, the inequality (4.3) can be written as

$$b_s(X) \leq a_s(a_t(X)) \qquad (s \leq t). \tag{4.4}$$

This is a natural strengthening of the standard bid-ask price inequality: the bid price at time s of the payoff X should be bounded above not only by the ask price of the payoff X at time s, but also by the ask price of any contract that allows locking in the payoff X at a later time t. By applying the inequality (4.3) to $-X$ instead of X, one obtains another strengthened form of the bid-ask price inequality:

$$b_s(b_t(X)) \leq a_s(X) \qquad (s \leq t) \tag{4.5}$$

The stronger inequality $b_s(b_t(X)) \leq b_s(X)$ is in fact equivalent to acceptance consistency, as shown below.

LEMMA 4.2 *A coherent dynamic evaluation ϕ is acceptance consistent if and only if, for all $X \in L^\infty$, we have*

$$\phi_s(X) \geq \phi_s(\phi_t(X)). \tag{4.6}$$

PROOF It is clear that the condition implies acceptance consistency. Conversely, given that ϕ is acceptance consistent, the fact that $\phi_t(X - \phi_t(X)) = 0$ implies that $\phi_s(X - \phi_t(X)) \geq 0$. Therefore, we can write

$$\phi_s(X) \geq \phi_s(\phi_t(X)) + \phi_s(X - \phi_t(X)) \geq \phi_s(\phi_t(X)). \tag{4.7}$$

This completes the proof. □

Consider now first the situation in which the bid-ask spread is zero for all assets, as is a standard assumption in a large part of the asset pricing literature. When the property $\phi_t(-X) = -\phi_t(X)$ is satisfied for all X, the inequality (2.4) turns into an equality, and consequently a coherent conditional evaluation generates a zero bid-ask spread if and only if it is linear. For families of linear coherent conditional evaluations, the notions of acceptance consistency, rejection consistency, sequential consistency, and strong consistency are all the same, as is readily verified on the basis of Lemma 4.2. For such dynamic evaluations, absence of arbitrage can be characterized as follows.

Proposition 4.3 *Consider a market in which prices at time t of are given by a family $\phi = (\phi_t)_{t \in \mathcal{T}}$ of linear coherent conditional evaluations. Such a market is free of arbitrage if and only if the family ϕ is sensitive and strongly consistent.*

Proof Suppose that the market is arbitrage-free and take $X \in L^\infty$. Given that $\phi_t(X) = a_t(X) = b_t(X)$ for all t, it follows from Lemma 4.1 that

$$\phi_s(\phi_t(X)) \leq \phi_s(X) \leq \phi_s(\phi_t(X))$$

for all s and t with $s \leq t$. In other words, we have $\phi_s(X) = \phi_s(\phi_t(X))$ so that strong consistency holds. To show sensitivity, suppose there exist $X \in L^\infty$ and $t \in \mathcal{T}$ such that $X \leq 0$ and $\phi_t(X) = 0$. Define $\Delta X_t = -X$, and note that $\phi_t(-\Delta X_t) = 0$ so that the condition in (4.2) is satisfied by taking $X_T = \Delta X_t = -X$. By absence of arbitrage, it follows that $X_T = 0$ and hence $X = 0$.

Next, assume that the family ϕ is sensitive and strongly consistent. Take X_T of the form (4.2) and suppose that $X_T \geq 0$; we need to show that $X_T = 0$. By linearity and strong consistency, we have

$$\phi_0(X_T) = \sum_{j=0}^{N} \phi_0(\Delta X_{t_j}) = \sum_{j=0}^{N} \phi_0(\phi_{t_j}(\Delta X_{t_j})) \leq 0.$$

On the other hand, we also have $\phi_0(X_T) \geq 0$ because $X_T \geq 0$, so that in fact $\phi_0(X_T) = 0$. The conclusion $X_T = 0$ now follows from linearity and sensitivity. □

The usual formulation of necessary and sufficient conditions for absence of arbitrage is of course in terms of an equivalent martingale measure. This formulation can be related to the proposition above by means of a representation theorem, similar to the representation theorems for strongly consistent families of coherent conditional evaluations as given for instance in [31].

When we now turn to situations in which the bid-ask spread is nonzero in general, the various notions of consistency are no longer equivalent. The following proposition shows that sensitivity and acceptance consistency are already sufficient to prevent arbitrage in the sense defined above.

Proposition 4.4 *Consider a market in which prices are determined by a coherent dynamic evaluation ϕ. If the family ϕ is sensitive and acceptance consistent, then the market is free of arbitrage.*

Proof Suppose that X_T is given by (4.2) with $X_0 = 0$, and that $X_T \geq 0$. By acceptance consistency, we have

$$\phi_0(-X_T) = \phi_0\Big(-\sum_{j=0}^{N} \Delta X_{t_j}\Big) \geq \sum_{j=0}^{N} \phi_0(-\Delta X_{t_j}) \geq 0.$$

From $X_T \geq 0$ we have $\phi_0(-X_T) \leq 0$, so that in fact $\phi_0(X_T) = 0$. It now follows from the assumed sensitivity that $X_T = 0$. □

The proposition shows that, from an arbitrage point of view, there is no need to impose the requirement of strong consistency. As was shown in Lemma 4.2, for coherent dynamic evaluations ϕ_t the inequality $\phi_s(X) \geq \phi_s(\phi_t(X))$ is implied by (in fact even equivalent to) acceptance consistency, and this in turn is implied both by conditional and by sequential consistency. An equivalent statement in terms of ask prices is

$$a_s(X) \leq a_s(a_t(X)) \qquad (s \leq t). \tag{4.8}$$

Strict inequality in the above means that it is less expensive at time s to buy the contract X directly than to buy at at time s a contract that gives the holder the opportunity to buy the contract X at the later time t. The difference of the right hand side and the left hand side can be viewed as a *postponement premium*, or vice versa as an *early decision discount*. The existence of such a premium/discount can be looked at as a form of illiquidity and does not give rise to arbitrage opportunities. Strong consistency comes down to replacing the inequality in (4.8) by an equality, or in other words to assuming that the postponement premium is zero.

5 Uniqueness of updating

It is a well known fact, recalled in Prop. 2.2 above, that any given conditional evaluation can be viewed as the conditional capital requirement corresponding to its acceptance set. The construction of the $\mathcal{F}_t$-refinement, 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 $\mathcal{F}_s$-conditional evaluation. Namely, given an $\mathcal{F}_s$-conditional evaluation ϕ_s, take its acceptance set $\mathcal{A}(\phi_s)$, construct the $\mathcal{F}_t$-refinement $(\mathcal{A}(\phi_s))^t$, and define the $\mathcal{F}_t$-update of ϕ_s as the conditional capital requirement that corresponds to $(\mathcal{A}(\phi_s))^t$. 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 $\mathcal{F}_t$-refinement of the acceptance set of the conditional evaluation at the earlier time or, more generally, the higher level of aggregation. Since as we have seen above conditional consistency is implied, under suitable sensitivity assumptions, by strong consistency as well as by sequential consistency, this implies that we obtain a criterion for a given conditional evaluation to be a member of a strongly or sequentially consistent family of conditional evaluations. Namely, it suffices to check whether the update that was described above is strongly or sequentially consistent.

To carry out this program, we first of all need to check whether the set $(\mathcal{A}(\phi_s))^t$ 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 5.1 *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 $\mathcal{F}_t$-refinement $\mathcal{S}^t$ 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). □

COROLLARY 5.2 *If the $\mathcal{F}_s$-conditional evaluation ϕ_s is sensitive, then the $\mathcal{F}_t$-refinement $(\mathcal{A}(\phi_s))^t$ of its acceptance set $\mathcal{A}(\phi_s)$ satisfies the basic properties of Prop. 2.2, namely acceptance of zero, solidness, and conditional nonnegativity.*

PROOF The acceptance set $\mathcal{A}(\phi_s)$ contains 0, is solid, and satisfies the property of negative cone exclusion by the assumption that ϕ_s is sensitive. The statement therefore follows from the proposition above. □

As a consequence, the following definition is justified.

DEFINITION 5.3 Let a sensitive $\mathcal{F}_s$-conditional evaluation ϕ_s be given, and let $t \geq s$. The *$\mathcal{F}_t$-refinement update* of ϕ_s is the $\mathcal{F}_t$-conditional evaluation ϕ_s^t defined by

$$\phi_s^t(X) = \operatorname{ess\,sup}\{Y \in L_t^\infty \mid \phi_s(1_F(X - Y)) \geq 0 \text{ for all } F \in \mathcal{F}_t\}. \tag{5.1}$$

The uniqueness of updating that was already mentioned above can now be stated more formally.

Theorem 5.4 *Let ϕ_s be a sensitive $\mathcal{F}_s$-conditional evaluation, and suppose that ϕ_t is a conditionally consistent update of ϕ_s. Then $\phi_t = \phi_s^t$, where ϕ_s^t is the refinement update as defined in (5.1).*

Proof Conditional consistency means that $\mathcal{A}_t = \mathcal{A}_s^t$. We know, as recalled in Prop. 2.2, that a conditional evaluation is the conditional capital requirement of its acceptance set. Therefore ϕ_t must be equal to the conditional capital requirement of $\mathcal{A}_s^t$, which by definition is ϕ_s^t. □

In view of Prop. 3.6, the statement above can be extended to uniqueness of sequential and of strong updates. Uniqueness of *strongly* 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 [35, Cor. 4.1]. The following theorem states the consequences of uniqueness of updating for existence of updates of a particular type.

Theorem 5.5 *Let ϕ_s be a sensitive $\mathcal{F}_s$-conditional evaluation, and let $t \geq s$. Then*

(i) ϕ_s allows a conditionally consistent $\mathcal{F}_t$-update if and only if the refinement update ϕ_s^t is a conditionally consistent update of ϕ_s

(ii) ϕ_s allows a strongly consistent $\mathcal{F}_t$-update if and only if the refinement update ϕ_s^t is a strongly consistent update of ϕ_s.

Moreover, if ϕ_s is strongly sensitive, then ϕ_s allows a sequentially consistent $\mathcal{F}_t$-update if and only if ϕ_s^t is a sequentially consistent update of $\mathcal{F}_s$.

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 an evaluation that does not allow a conditionally consistent update.

Example 5.6 Consider a nonrecombining two-step binomial tree. The sample space is $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, P is the uniform measure, $\mathcal{F} = \mathcal{F}_2 = 2^\Omega$, and $\mathcal{F}_1$ is generated by $\{\omega_1, \omega_2\}$. The space L^∞ can in this case be identified with $\mathbb{R}^4$. Define a concave unconditional evaluation $\phi_0 : \mathbb{R}^4 \to \mathbb{R}$ 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 evaluation 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 evaluation ϕ_0 does not admit a conditionally consistent update, and neither does it allow a sequentially consistent or a strongly 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 evaluation.

Consider now a *dynamic evaluation* 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 $t \in \mathcal{T}$, ϕ_t is an $\mathcal{F}_t$-conditional evaluation. It would be reasonable to speak of a *strongly consistent dynamic evaluation* if, for each $s, t \in \mathcal{T}$ with $s \leq t$, ϕ_t is a strongly consistent update of ϕ_s, and likewise for other notions of consistency. The conditions in Thm. 5.5 provide criteria for an initial evaluation 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 [27, Def. 2.1.2, Prop. 2.1.6] (cf. also [33, 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 consistency. This is a consequence of the following proposition.

PROPOSITION 5.7 *Let ϕ_s, ϕ_t, and ϕ_u be conditional evaluations, 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 "sequentially consistent" or by "strongly consistent".*

The proof of the proposition is given in the Appendix.

6 Membership of consistent families

It was shown above that the refinement update provides the unique candidate for consistent updating. As shown in Example 5.6, the refinement update need not be conditionally consistent, let alone sequentially or strongly consistent. This means that, given an evaluation ϕ_0 at the highest level of aggregation, it is not in general true that this evaluation is a member of a consistent family of evaluations, 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 evaluation 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. 5.4, 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 (3.3). In fact it is enough to verify acceptance consistency, as shown in the following proposition.

PROPOSITION 6.1 *Let ϕ_t be the $\mathcal{F}_t$-refinement update of a given $\mathcal{F}_s$-conditional evaluation ϕ_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 (3.3) follows from the definition of the refinement update. Moreover (as noted by Tutsch [33, 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 $1_F X$. □

In other words, the proposition states that if the refinement update of a given conditional evaluation is not conditionally consistent, then it is in fact not even acceptance consistent. This is the situation we encountered in Example 5.6.

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 $\mathcal{F}_t$-refinement of the acceptance set $\mathcal{A}_s$ to qualify as the acceptance set of an $\mathcal{F}_t$-conditional evaluation. 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 6.2 *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 $\mathcal{F}_t$-refinement of the acceptance set of ϕ is $\mathcal{F}_t$-closed.*

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

PROPOSITION 6.3 *A sensitive $\mathcal{F}_s$-conditional evaluation ϕ_s that is continuous from above admits a conditionally consistent $\mathcal{F}_t$-update if and only if it has the following property:*

$$[\,\exists G \in \mathcal{F}_t : \ \forall F \in \mathcal{F}_t : \ \phi_s(1_{F\cap G}X) \geq 0, \ \phi_s(1_{F\cap G^c}X) \geq 0\,] \ \Rightarrow \ \phi_s(X) \geq 0 \quad (X \in L^\infty). \tag{6.1}$$

Proof As already noted above, under the stated conditions the evaluation ϕ_s admits a conditionally consistent $\mathcal{F}_t$-update if and only if the $\mathcal{F}_t$-refinement $\mathcal{A}_s^t$ 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 $\mathcal{F}_t$-refinements, so that it only remains to show that the property (6.1) 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

$$[\,\exists G \in \mathcal{F}_t :\ \forall F \in \mathcal{F}_t :\ \phi_s(1_{F\cap G}X) \geq 0,\ \phi_s(1_{F\cap G^c}X) \geq 0\,] \ \Rightarrow\ \forall F \in \mathcal{F}_t :\ \phi_s(1_F X) \geq 0. \tag{6.2}$$

This property obviously implies (6.1). Conversely, suppose that (6.1) holds, and let $X \in L^\infty$ be such that the condition on the left hand side of (6.2) is fulfilled. Take $H \in \mathcal{F}_t$. The condition in (6.2) is then fulfilled also for $1_H X$ instead of X, and since this premise is the same as in (6.1) it follows that $\phi_s(1_H X) \geq 0$. Therefore the conclusion of (6.2) holds. Consequently (6.2) is equivalent to (6.1). □

We will refer to (6.1) 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

$$[\,\exists G \in \mathcal{F}_t : \phi_s(1_G X) \geq 0,\ \phi_s(1_{G^c}X) \geq 0\,] \ \Rightarrow\ \phi_s(X) \geq 0 \quad (X \in L^\infty). \tag{6.3}$$

Indeed, the premise of (6.1) implies the one in (6.3).

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 (6.4) that is used below may be called *generic complementarity.*

Lemma 6.4 *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:*

$$X, Y \in \mathcal{S},\ P\big(\{X = 0\} \cup \{Y = 0\}\big) = 1 \ \Rightarrow\ X + Y \in \mathcal{S}. \tag{6.4}$$

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 [29, 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 evaluations that are not of the coherent ("multiple-prior") type.

PROPOSITION 6.5 *Let $\mathcal{Q}$ be a collection of probability measures that are all absolutely continuous with respect to the reference measure P. Define*

$$\phi_t(X) = \operatorname*{ess\,inf}_{Q\in\mathcal{Q}} E_t^Q X \qquad (X \in L^\infty,\ t \in \mathcal{T}). \tag{6.5}$$

Assume that ϕ_0 is sensitive. The dynamic evaluation $(\phi_t)_{t\in\mathcal{T}}$ 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 9.4 in the Appendix, we can conclude from this that $E_t^Q X \geq 0$. □

A sufficient but not necessary condition for the sensitivity requirement in the proposition to be satisfied is that the collection $\mathcal{Q}$ contains at least one measure that is equivalent to P.

Necessary and sufficient conditions for the existence of sequentially consistent updates were given for the class of law-invariant concave evaluations by Weber [35, Thm. 4.3, 4.4]. The strongly consistent families within the same class have been fully described by Kupper and Schachermayer [24]. Alternative characterizations (not requiring law invariance) are provided in the following two propositions.

PROPOSITION 6.6 *A strongly sensitive $\mathcal{F}_s$-conditional evaluation 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*

$$\phi_s(1_F(X - C_t)) = 0 \quad (F \in \mathcal{F}_t). \tag{6.6}$$

PROPOSITION 6.7 *A strongly sensitive $\mathcal{F}_s$-conditional evaluation ϕ_s admits a strongly 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*

$$\phi_s(1_F X) = \phi_s(1_F C_t) \quad (F \in \mathcal{F}_t). \tag{6.7}$$

The proofs of both propositions are in the Appendix.

7 Families of compound risk measures

A standard example of a strongly consistent dynamic evaluation is the family of entropic conditional evaluations that is defined, for $t \in \mathcal{T} = \{0, 1, \dots, T\}$, by

$$\phi_t^\beta(X) = \frac{-1}{\beta} \log E_t^P \exp(-\beta X)$$

where P is a given measure and β is a constant; the sign convention in the expression above is such that the evaluations ϕ_t^β are concave when β is positive. When $\beta = 0$, the right hand side is replaced by $E_t^P X$. The entropic dynamic evaluation has only one parameter β which controls the amount of emphasis on adverse outcomes both for short and for long horizons. A larger class that allows more flexibility in this respect may be constructed as follows. First, let aggregation from level $t+1$ to level t be carried out by an entropic evaluation with parameter β_t:

$$\bar{\phi}_t(X) = -\frac{1}{\beta_t} \log E_t^P \exp(-\beta_t X) \qquad (X \in L_{t+1}^\infty) \tag{7.1}$$

The mappings $\bar{\phi}_t : L_{t+1}^\infty \to L_t^\infty$ that are obtained in this way can subsequently be pieced together to form conditional evaluations:

$$\phi_T^\beta(X) = X \tag{7.2a}$$

$$\phi_t^\beta(X) = \bar{\phi}_t(\phi_{t+1}^\beta(X)) \tag{7.2b}$$

where $\beta := (\beta_0, \beta_1, \dots, \beta_{T-1})$. In this way one obtains a collection of dynamic evaluations $\phi^\beta = (\phi_0^\beta, \dots, \phi_{T-1}^\beta)$, parametrized by the constants β_t which relate to short-horizon evaluations. One can think of the β_t's as describing a particular amount of risk tolerance that is applied between times t and $t+1$. As a next step, one can the introduce a "budget" for risk tolerance which may be spent over a number of time steps. This idea can be implemented by defining *compound* conditional evaluations as follows:

$$\Phi_t^\beta(X) = \operatorname{ess\,inf}\{\phi_t^\beta(X) \mid \textstyle\sum_{s=t}^{T-1} \beta_s \le b(t)\} \tag{7.3}$$

where $b(t)$ is a given deterministic "profile function". The construction might be generalized further, for instance by allowing the β_t's to be $\mathcal{F}_t$-measurable functions rather than constants.

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

$$\hat{\phi}_t(X) = \operatorname*{ess\,inf}_{\phi \in \Phi} \phi_t(X) \qquad (X \in L^\infty). \tag{7.4}$$

It is easily verified that the mappings $\hat{\phi}_t$ $(0 \le t \le T)$ are indeed conditional evaluations. The evaluation that is obtained in this way is an example of a *compound* evaluation. More generally one might aggregate by using any (unconditional) evaluation functional, instead of the infimum.

The dynamic evaluations obtained from (7.4) are in general not strongly consistent, even if all dynamic evaluations in the collection Φ do have this property. Acceptance consistency holds, however; this follows from the lemma below.

LEMMA 7.1 *Let Φ be a collection of dynamic evaluations $\phi = (\phi_t)_{t\in\mathcal{T}}$. Define a dynamic evaluation $\hat{\phi} = (\hat{\phi}_t)_{t\in\mathcal{T}}$ by (7.4). If all dynamic evaluations $\phi \in \Phi$ are strongly consistent, then the dynamic evaluation $\hat{\phi}$ 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. We work in a discrete-time setting and construct strongly consistent conditional evaluations in a step-by-step manner, as in the example above. To describe this in general terms, we follow Cheridito and Kupper [11]. Specifically, suppose that, for $t = 0, \ldots, T-1$, one-step evaluations $\bar{\phi}_t : L^\infty_{t+1} \to L^\infty_t$ are given, and define a corresponding sequence of conditional evaluations by the backward recursion (7.2). 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. We consider the situation in which the one-step evaluations $\bar{\phi}_t$, called "generators" in [11], are concave and continuous from above so that they can be represented in terms of penalty functions [15, Thm. 4.16]. First define, for $t = 1, \ldots, T$, sets of one-step densities by

$$\mathcal{D}_t = \{\xi \in L^1_t \mid \xi \geq 0,\ E^P_{t-1}\xi = 1\}.$$

For a given probability measure $Q \ll P$, write $M^Q_t = E^P_t(dQ/dP)$ and define $\xi^Q_t \in \mathcal{D}_t$ by $\xi^Q_t = M^Q_t/M^Q_{t-1}$ on $\{M^Q_{t-1} > 0\}$ and $\xi^Q_t = 1$ otherwise. Conversely, given any sequence of one-step densities $(\xi_1, \ldots, \xi_T)$, a probability measure $Q \ll P$ that satisfies $\xi^Q_t = \xi_t$ for all t is obtained from

$$E^Q X = E^P[\xi_1 \cdots \xi_T X] \qquad (X \in L^\infty).$$

A mapping $\bar{\theta}_t$ from $\mathcal{D}_{t+1}$ to the set of $\mathcal{F}_t$-measurable functions is said to be a *one-step penalty function* if it satisfies the following properties [11, Def. 3.3]:

$$\operatorname*{ess\,inf}_{\xi\in\mathcal{D}_{t+1}} \bar{\theta}_t(\xi) = 0 \tag{7.5}$$

$$\bar{\theta}_t(1_F\xi + 1_{F^c}\xi') = 1_F\bar{\theta}_t(\xi) + 1_{F^c}\bar{\theta}_t(\xi') \quad (\xi, \xi' \in \mathcal{D}_{t+1},\ F \in \mathcal{F}_t). \tag{7.6}$$

For $Q \ll P$, write $\bar{\theta}_t(Q) = \bar{\theta}_t(\xi^Q_{t+1})$. A *dynamic penalty function* is a sequence $(\bar{\theta}_t)_{t=0,\dots,T-1}$ of one-step penalty functions. Generators associated to a dynamic penalty function can now be defined by

$$\bar{\phi}_t(X) = \operatorname*{ess\,inf}_{Q \ll P} \big(E^Q_t X + \bar{\theta}_t(Q)\big).$$

It is shown by Cheridito and Kupper [11, Thm. 3.4] that, when the generators are defined in this way, the evaluations defined by (7.2) can be represented as follows:

$$\phi_t(X) = \operatorname*{ess\,inf}_{Q \ll P} E^Q_t X + \theta_t(Q) \tag{7.7}$$

with

$$\theta_t(Q) = E^Q_t \sum_{u=t}^{T-1} \bar{\theta}_u(Q). \tag{7.8}$$

In general, suppose that a collection $\bar{\Theta}$ of dynamic penalty functions is given. This collection gives rise to a collection Φ of dynamic evaluations, each of which is strongly consistent by construction. Now form a new dynamic evaluation by the formula (7.4). If we let Θ denote the class of sequences of functions obtained from $\bar{\Theta}$ by (7.8), the dynamic evaluation that is obtained in this way can be written as

$$\hat{\phi}_t(X) = \operatorname*{ess\,inf}_{\theta \in \Theta} \operatorname*{ess\,inf}_{Q \ll P} E^Q_t X + \theta_t(Q). \tag{7.9}$$

Consider now in particular the situation in which the measure P is taken as a "central" measure and the penalty function takes the form of a distance with respect to this measure. For instance entropic risk measures can be described as such; the Gini index provides another example. Then the following property is satisfied:

$$\bar{\theta}_t(P) = 0 \text{ for all } \bar{\theta} \in \bar{\Theta},\ t = 0, \dots, T-1. \tag{7.10}$$

Under this assumption, we can prove conditional consistency.

Proposition 7.2 *Let the dynamic evaluation $\hat{\phi}$ be defined by (7.4), where the family Φ of dynamic evaluations is obtained from a family $\bar{\Theta}$ of dynamic penalty functions, and assume that (7.10) is satisfied. Then $\hat{\phi}$ is conditionally consistent.*

Proof Take s and t with $0 \le s < t \le T$, and let X be an element of L^∞. We need to show that $\hat{\phi}_t(X) \ge 0$ if and only if $\hat{\phi}_s(1_F X) \ge 0$ for all $F \in \mathcal{F}_t$. Suppose first that $\hat{\phi}_t(X) \ge 0$. For $F \in \mathcal{F}_t$, we then have, by Lemma 7.1, $\hat{\phi}_s(1_F X) \ge \hat{\phi}_s(\hat{\phi}_t(1_F X)) = \hat{\phi}_s(1_F \hat{\phi}_t(X)) \ge 0$.

Next, assume that $\phi_s(1_F X) \ge 0$ for all $F \in \mathcal{F}_t$, and suppose that $\hat{\phi}_t(X) \ngeq 0$ for some $X \in L^\infty$. Then there must exist a function $\theta \in \Theta$, a measure $Q \ll P$, an $F \in \mathcal{F}_t$ with $P(F) > 0$ and an $\varepsilon > 0$ such that

$$1_F\big(E^Q_t X + \theta_t(Q)\big) \le -\varepsilon 1_F. \tag{7.11}$$

Define a probability measure Q' by

$$\xi_u^{Q'} = \begin{cases} 1 & (u \le t) \\ 1_F \xi_u^Q + 1_{F^c} & (u > t). \end{cases}$$

Under this definition, we have

$$\begin{aligned} E_t^{Q'}(Y) &= E_t^P\big((1_F\xi_{t+1}^Q + 1_{F^c}) \cdots (1_F\xi_T^Q + 1_{F^c})Y\big) = E_t^P\big((1_F\xi_{t+1}^Q \cdots \xi_T^Q + 1_{F^c})Y\big) \\ &= 1_F E_t^Q Y + 1_{F^c} E_t^P Y \end{aligned}$$

for all $Y \in L^\infty$. In particular, $1_F E_t^{Q'} Y = 1_F E_t^Q Y$ for all Y. Therefore we can write, using the assumption (7.10),

$$\begin{aligned} 1_F\theta_t(Q) &= 1_F E_t^Q \sum_{u=t}^{T-1} \bar\theta_u(Q) = E_t^{Q'} \sum_{u=t}^{T-1} 1_F\bar\theta_u(Q) = E_t^{Q'} \sum_{u=t}^{T-1} \big(1_F\bar\theta_u(\xi_{u+1}^Q) + 1_{F^c}\bar\theta_u(1)\big) \\ &= E_t^{Q'} \sum_{u=t}^{T-1} \big(\bar\theta_u(1_F\xi_{u+1}^Q + 1_{F^c})\big) = E_t^{Q'} \sum_{u=t}^{T-1} \bar\theta_u(\xi_{u+1}^{Q'}) \\ &= \theta_t(Q'). \end{aligned}$$

We can now rewrite (7.11) as

$$E_t^{Q'}(1_F X) + \theta_t(Q') \le -\varepsilon 1_F. \tag{7.12}$$

Note that

$$\theta_s(Q') = E_s^{Q'} \sum_{u=s}^{T-1} \bar\theta_u(Q') = E_s^{Q'} \sum_{u=s}^{T-1} \bar\theta_u(\xi_{u+1}^{Q'}) = E_s^{Q'} \sum_{u=t}^{T-1} \bar\theta_u(Q') = E_s^P \theta_t(Q').$$

Therefore, applying the conditional expectation operator E_s^P to both sides of (7.12), we obtain

$$E_s^{Q'}(1_F X) + \theta_s(Q') \le E_s^P(-\varepsilon 1_F). \tag{7.13}$$

Since we assumed that $\phi_s(1_F X) \ge 0$ for all $F \in \mathcal{F}_t$, the left hand side is nonnegative. This implies that $-\varepsilon P(F) = E^P E_s^P(-\varepsilon 1_F) \ge 0$, so that we have reached a contradiction. □

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

$$\bar\theta, \bar\theta' \in \bar\Theta,\ F \in \mathcal{F}_t \Rightarrow \exists \bar\theta'' \in \bar\Theta \text{ s.t. } \bar\theta_u = 1_F\bar\theta_u + 1_{F^c}\bar\theta'_u \quad (0 \le t \le T,\ u \ge t). \tag{7.14}$$

Proposition 7.3 *Under the assumption (7.14) and the assumptions of Prop. 7.2, the dynamic evaluation defined by (7.4) 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 $\hat\phi_s(X) \ge 0$ follows from the assumption $\hat\phi_t(X) = 0$ by Lemma 7.1. It remains to prove the reverse inequality. To this end we show that the family $\mathcal{E} := \{E_t^Q X + \theta_t(Q) \,|\, Q \ll P,\ \theta \in \Theta\}$ is directed downwards. Let two elements of this family be given by Y and Y', corresponding to (Q, θ) and (Q', θ') respectively, and write $F = \{Y \le Y'\} \in \mathcal{F}_t$. Define the measure Q'' by

$$\xi_u^{Q''} = \begin{cases} 1 & (u \le t) \\ 1_F \xi_u(Q) + 1_{F^c} \xi_u(Q') & (u > t). \end{cases}$$

Let $\bar{\theta}'' \in \bar{\Theta}$ be such that $\bar{\theta}_u = 1_F \bar{\theta}_u + 1_{F^c} \bar{\theta}'_u$ for $u \ge t$, and write $Y'' = E_t^{Q''} X + \theta''_t(Q'')$. We have $\min(Y, Y') = Y'' \in \mathcal{E}$, so that indeed the family $\mathcal{E}$ is directed downwards. It now follows (see for instance [15, Thm. A.33]) that there exists a sequence $(Q_k, \theta^k)_{k \in \mathbb{N}}$ such that

$$E_t^{Q_k} X + \theta_t^k(Q_k) \searrow 0. \tag{7.15}$$

We may assume that the Q_k's have been chosen such that $\xi_u^{Q_k} = 1$ for $u \le t$. We then have $E_s^{Q_k} X = E_s^P E_t^{Q_k} X$ and $\theta_s(Q_k) = E_s^P \theta_t(Q_k)$. The monotone convergence theorem applied to (7.15) implies that $\lim_{k \to \infty} E_s^{Q_k} X + \theta_s^k(Q_k) = 0$, and consequently we must have $\hat{\phi}_s(X) = 0$. $\square$

8 Conclusions

The construction of dynamic risk measures that combine time consistency with reasonable levels of prudence across different time horizons remains a challenging task. In this paper we have shown that, even under very weak interpretations of the notion of time consistency, the choice of the initial risk measure already fully determines the family of evaluations to which it belongs. By considering different notions of time consistency, we obtain a categorization of risk measures in measures that allow conditionally consistent updating, measures that allow sequentially consistent updating, and measures that allow strongly consistent updating. We have provided characterizations of these properties, 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. Conditions for weak consistency in terms of dual representations are provided in [30].

9 Appendix

9.1 Auxiliary results

We prove a few auxiliary results. The argument in the proof of the first lemma is similar to the reasoning in the proof of Thm. 4.33 in [15].

LEMMA 9.1 *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 $\phi(X) \geq 0$.*

PROOF Define $Y_n = \operatorname{ess\,sup}_{m\geq n} X_m$; then $Y_n \searrow X$ so that $\phi(X) = \lim_{n\to\infty} \phi(Y_n)$. By monotonicity and normalization, we have $\phi(Y_n) \geq \phi(X_n) \geq 0$ for all n, so that $\lim_{n\to\infty} \phi(Y_n) \geq 0$ and the stated result follows. □

LEMMA 9.2 *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 $X_t \not\geq Y_t$, 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 9.3 *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 9.4 *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 9.2 with $Y_t = 0$, noting that $\phi(Y_t) = \phi(0) = 0$ and replacing strong sensitivity with sensitivity. □

9.2 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 evaluation 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 [13] that the real-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 a conditional evaluation ϕ_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].

9.3 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 consistency. It remains to prove that the reverse inclusion holds under each of the three conditions mentioned. That is, we need to show the implication from right to left in (3.3). 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 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:

$$0 \leq \phi_s(1_F X) \leq \phi_s(1_F X - 1_F\phi_t(X)) = \phi_s(1_F X - \phi_t(1_F X)) = 0. \tag{9.1}$$

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 (9.1) 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 $1_F \phi_t(X) = 0$.

9.4 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 $\mathcal{F}_t$-refinement 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 $\phi_t(X) \geq Z$.

9.5 Proof of Lemma 6.2

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 9.1, 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 9.1 again, we conclude that $\phi(1_F X) \geq 0$. Since $F \in \mathcal{F}_t$ was arbitrary, it follows that $X \in \mathcal{S}^t$.

9.6 Proof of Prop. 5.7

The property expressed in Prop. 5.7 for sequentially or strongly 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 $\phi_u(X) \geq 0$.

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 $\phi_t(X) \ngeq 0$, 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 evaluation 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 consistency

Assume that ϕ_u is a strongly 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 consistent update of ϕ_s. Conversely, assume now that ϕ_t and ϕ_u are strongly 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. 9.3 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 consistency.

Lemma 9.5 *A sequentially consistent update of a strongly sensitive conditional evaluation 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 ϕ_s and ϕ_t be conditional evaluations with respect to $\mathcal{F}_s$ and $\mathcal{F}_t$ respectively. 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 $X = Y$. $\square$

9.7 Proof of Prop. 6.6

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 ϕ_s is an $\mathcal{F}_s$-conditional evaluation that has a conditionally consistent $\mathcal{F}_t$-update, say ϕ_t, and that for each $X \in L^\infty$ there exists $C_t \in L_t^\infty$ such that (6.6) 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 (6.6) holds. By conditional consistency, the condition (6.6) 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 (6.6), 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

$$\operatorname{ess\,sup}\{Y_t \in L_t^\infty \mid \phi_s(1_F(X - C_t - Y_t)) \geq 0 \text{ for all } F \in \mathcal{F}_t\} \leq 0. \tag{9.2}$$

The conditional evaluation ϕ_t must be equal to the refinement update of ϕ_s, by Thm. 5.4. In view of the expression given for the refinement update in (5.1), it follows from (9.2) that $\phi_t(X - C_t) \leq 0$. Therefore, we obtain the inequality $\phi_t(X) \leq C_t$, and the proof is complete.

9.8 Proof of Prop. 6.7

If ϕ_s admits a strongly 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 (6.7) holds. It follows from Cor. 9.3 in the Appendix 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

$$\phi_s(1_F X) = \phi_s(1_F \psi(X)) \quad (F \in \mathcal{F}_t). \tag{9.3}$$

If we can show that the mapping ψ is an $\mathcal{F}_t$-conditional evaluation, then strong consistency follows from (9.3) and the proof will be complete.

In order to prove that ψ is an $\mathcal{F}_t$-conditional evaluation, it suffices [9, Rem. 3.4] to prove that ψ is normalized and monotonic, and that it satisfies the local property as well as real 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 9.2 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 real 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 real 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 (9.3), 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] P. Cheridito and M. Kupper. Composition of time-consistent dynamic monetary risk measures in discrete time. NCCR FINRISK Working Paper 409, ETH Zürich, 2006.

[12] 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.

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

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

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

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

[17] 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.

[18] E. Hanany and P. Klibanoff. Updating preferences with multiple priors. *Theoretical Economics*, 2:261–298, 2007.

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

[20] E. Jouini and H. Kallal. Martingales and arbitrage in securities markets with transaction costs. *Journal of Economic Theory*, 66:178–197, 1995.

[21] Y. Kabanov. Hedging and liquidation under transaction costs in currency market. *Finance and Stochastics*, 3:237–248, 1999.

[22] Y. Kabanov and M. Safarian. *Markets with Transaction Costs. Mathematical Theory.* Springer, Berlin, 2009.

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

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

[25] M. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27:1622–1668, 1989.

[26] S. Peng. Filtration consistent nonlinear expectations and evaluations. In *3rd Workshop on Markov Processes and Related Topics,* Beijing, August 10–14, 2004.

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

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

[29] 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.

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

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

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

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

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

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