Valuations and dynamic convex risk measures.

A. Jobert*
University of Cambridge

L.C.G. Rogers†
University of Cambridge

August 2, 2005

Abstract

This paper approaches the definition and properties of dynamic convex risk measures through the notion of a family of concave valuation operators satisfying certain simple and credible axioms. Exploring these in the simplest context of a finite time set and finite sample space, we find natural risk-transfer and time-consistency properties for a firm seeking to spread its risk across a group of subsidiaries.

1 Introduction.

The growing literature of risk measurement considers mainly[1] single-period risk measurement, where one attempts to 'measure' at time zero the risk involved in undertaking

*Statistical Laboratory, University of Cambridge, Wilberforce Road, Cambridge CB30WB, UK; A.Jobert@statslab.cam.ac.uk

†Statistical Laboratory, University of Cambridge, Wilberforce Road, Cambridge CB30WB, UK; L.C.G.Rogers@statslab.cam.ac.uk The authors thank participants at the Isaac Newton Institute programme Developments in Quantitative Finance 2005 for helpful discussions and comments. We thank particularly Philippe Artzner, Alexander Cherny, Phil Dybvig, Marco Frittelli, Lars Hansen, and Jose Scheinkman. Any remaining errors are of course the responsibility of the authors.

[1]See [2, 16, 11, 1, 17, 10, 9, 23, 12, 18, 20, 7] for one-period risk-measurement, [3] for the multiperiod extension of [2], [28, 5, 8, 25] for a particular class of dynamic risk measures (the multiperiod behaviour addressed by the latter papers is somehow less general than the one considered by [3] and will be commented later), [19, 21, 26, 27] for further dynamic risk measures, and [22, 14] for the related economics literature devoted to preference relations and Bayesian decision-making.

to receive some contingent claim X at time 1. In this literature, a set $\mathcal{A}$ of *acceptable* contingent claims is frequently taken to be the primitive object (as in [2], for example). Such a set gives rise naturally to a *risk measure* $\pi_{\mathcal{A}}$ via the definition

$$\pi_{\mathcal{A}}(X) = \sup\,[\, m \,|\, X - m \in \mathcal{A}\,]\,,$$

which is simply the greatest price at which it would be acceptable to buy the contingent claim X. Artzner et al. [2] define a *coherent utility function*[2] to be one which satisfies five axioms equivalent to

(CRM1) *concavity*: $\pi(\lambda X + (1-\lambda)Y) \geq \lambda\pi(X) + (1-\lambda)\pi(Y) \quad (0 \leq \lambda \leq 1)$;

(CRM2) *positive homogeneity*: if $\lambda \geq 0$, then $\pi(\lambda X) = \lambda\pi(X)$;

(CRM3) *monotonicity*: if $X \leq Y$, then $\pi(X) \leq \pi(Y)$;

(CRM4) *translation invariance*: if $m \in \mathbf{R}$, then $\pi(Y+m) = \pi(Y) + m$.

(CRM5) *relevance*: if $X \geq 0$, and $X \neq 0$ then $\pi(X) > 0$.

They go on to show that (under simplifying assumptions) any such risk measure is representable as[3]

$$\pi(X) = \inf_{Q \in \mathcal{Q}} E_Q[X], \tag{1}$$

where $\mathcal{Q}$ is some collection of probability measures[4].

The positive-homogeneity condition (CRM2) is arguably unnatural, and was removed by Föllmer & Schied [16] and by Frittelli & Gianin [18] who thereby introduced the notion of a *convex risk measure*. They show that a convex risk measure admits a representation as

$$\pi(X) = \inf_{Q \in \mathcal{Q}} \{\, E_Q[X] - \alpha(Q)\,\} \tag{2}$$

where α is a concave 'penalty' function on $\mathcal{Q}$. Clearly if $\alpha \equiv 0$, then we recover the representation of a coherent risk measure, but the notion of a convex risk measure is more general.

Of course, the usefulness of a single-period study should be judged by the extent to which it helps us to understand risk measurement in a multi-period setting; this has been well recognised for some time, and recently attempts have been made to achieve that extension. For example, Artzner, Delbaen, Eber, Heath and Ku [3] adapt the static coherent risk measure axioms to the product sample space $\{0, 1, \ldots, T\} \times \Omega$ and obtain a representation for a coherent risk measure of the form:

$$\pi(X) = \inf_{V \in \mathcal{V}} E_P\left[\sum_{t=0}^{T} X_t(V_t - V_{t-1})\right]$$

[2] In fact, the paper [2] works with coherent risk measures ρ, whose negatives $-\rho$ satisfy the given conditions. In that paper, coherent risk measures are related to the primitive concept, the acceptance set.

[3] The properties (CRM) appeared in an earlier paper of Gilboa & Schmeidler [22], in the context of Bayesian decision theory. This study was not concerned with risk measurement.

[4] Evidently, if π has the form (1) then it satisfies the properties (CRM1-4).

for one fixed probability measure P and a set $\mathcal{V}$ of positive increasing adapted processes. The argument of π is of course a *cash-flow* process.

Föllmer and Schied [16, 17], Cvitanič and Karatzas [8] and Nakano [25] also have several periods for portfolio construction and measure the risk of final wealths, which may depend on trajectories. However, they do not incorporate in their analysis what happens at intermediate dates. The trajectories have only an impact on the final values. Similarly, in Riedel [28], the risk-adjusted measurement of cumulative cash-flows involves only final values and not the whole trajectories. By contrast, as in [3], we not only look at final values but also at intermediate time points. Although Cheredito, Delbaen and Kupper [5] consider continuous-time discounted value processes, the convex risk measures treated in their paper are static as they only measure the risk of a discounted value process at the beginning of a given time period.

The major issue which arises in a multiperiod framework is the one of *dynamic consistency.* Although every set of probability measures generates a coherent risk measure in the static framework, only sets of probability measures consistent in an appropriate sense yield dynamic coherent risk measures. This consistency property of probability measures (or stability by "pasting") has been analysed by Epstein and Schneider [14] (building upon the atemporal multiple-priors model of Gilboa and Schmeidler [22] and using prior-by-prior Bayesian updating for "rectangular" sets of priors), Artzner et al. [3] (using change-of-measure martingales) and Riedel [28] (via Bayesian updating and a different kind of translation invariance property). It is often referred to as *multiplicative stability* [11]. After the first draft of this paper was complete, we learned of a preprint of Cheredito, Delbaen & Kupper [6] which develops an axiomatic framework quite similar to ours. The main aim there appears to be to explore the implications of the given setup for acceptance sets and coherent risk measures, relating to earlier work of the authors. Our own emphasis is quite different; we do not concern ourselves with acceptance sets, but rather wish to understand what consequences of the axiomatic setup can be developed. We will comment further on the relations between the two contributions as the occasions arise.

In this paper, we define and analyse[5] the notion of a *dynamic convex risk measure*, extending the dynamic coherent risk measure of [3], rather as Frittelli & Gianin [18] and Föllmer & Schied [16] extend [2] in the single-period context. However, we do not begin with the notion of a set of acceptable cash-flows, but start from a family of *valuation operators*, or, more briefly, *valuations.*

As a first hint of the usefulness of this general approach, we quote a simple result which is presumably well known (it certainly appears in Rogers [30], for example.) The idea is to write down certain natural axioms that *market valuation operators* should have, and to derive implications.

[5]We work in the technically simple setting of a finite time set, and a finite probability space Ω; this allows us to obtain the main ideas without being held up by technical issues.

Theorem 1 *In a filtered probability space* $(\Omega, (\mathcal{F}_t)_{t\geq 0}, P)$*, suppose that valuation operators* $(\pi_{tT})_{0\leq t\leq T}$

$$\pi_{st} : L^\infty(\mathcal{F}_t) \to L^\infty(\mathcal{F}_s) \qquad (0 \leq s \leq t).$$

satisfy the following four axioms:

(A1) Each π_{st} *is a bounded positive linear operator from* $L^\infty(\mathcal{F}_t)$ *to* $L^\infty(\mathcal{F}_s)$*;*

(A2) If $Y \in L^\infty(\mathcal{F}_t)$*,* $Y \geq 0$*, then*

$$\pi_{0t}(Y) = 0 \iff P(Y > 0) = 0.$$

*(*no arbitrage*)*

(A3) For $0 \leq s \leq t \leq u$*,* $Y \in L^\infty(\mathcal{F}_u)$*,* $X \in L^\infty(\mathcal{F}_t)$*,*

$$\pi_{su}(XY) = \pi_{st}(X\pi_{tu}(Y))$$

*(*dynamic consistency*)*

(A4) If $(Y_n) \in L^\infty(\mathcal{F}_t)$*,* $|Y_n| \leq 1$*,* $Y_n \uparrow Y$ *then* $\pi_{st}(Y_n) \uparrow \pi_{st}(Y)$ *(*continuity*)*

For simplicity, suppose also that $\mathcal{F}_0$ *is trivial. Then there exists a strictly positive process* $(\zeta_t)_{t\geq 0}$ *such that the valuation operators* π_{st} *can be expressed as*

$$\pi_{st}(Y) = \frac{E\big[\,\zeta_t Y \,\big|\, \mathcal{F}_s\,\big]}{\zeta_s} \quad (0 \leq s \leq t). \tag{3}$$

The proof of this result takes about a page, and is included in the appendix; nothing more sophisticated than standard facts about measure theory is required[6]. However, its importance is not to be underestimated; it is in fact a *poor man's Fundamental Theorem of Asset Pricing* (FTAP). Indeed, the conclusion of Theorem 1 is exactly what comes out of the FTAP, but its axiomatic starting point is different; in the usual FTAP we start from some suitably-formulated axiom of absence of arbitrage, and here we start from the axioms (A1)–(A4). Which of these two axiomatic starting points one should wish to assume is of course a matter of taste; in defence of the unconventional approach taken here, it is worth pointing out[7] that if we want to have the conclusion (3), then (A1)–(A4) must hold anyway!

Of the four axioms assumed in Theorem 1, the key one is the dynamic consistency axiom, (A3), as you will see from the proof; without this, we are able to prove that (3) holds if $s = 0$, but this is of course far too limited to be useful. Notice the interpretation of (A3); we can obtain X units of Y at time u in two ways, either by buying at time s the contingent claim XY, or by buying at time s the contingent claim which at time t will deliver X units of the time-t price $\pi_{tu}(Y)$ of Y, and (A3) says that these two should be valued the same at time s.

[6]Note however that the assumption that the valuation operators are defined on the *whole* of $L^\infty(\mathcal{F}_t)$ greatly simplifies the argument.

[7]It is also worth pointing out that it took years to find the correct formulation for the notion of absence of arbitrage!

Now Theorem 1 relates to *market* valuations, where linearity in the contingent claim being priced is a reasonable assumption; if we want to buy X and Y, the price will be the price of X plus the price of Y. However, when it comes to risk measurement, what the valuation operator is doing is to tell us how much capital a given firm should set aside to allow it to accept a named cash balance. Linearity now would *not* be a property that we want (we might require a positive premium both to cover a cash balance C and to cover $-C$, but we would not require a positive premium to cover the sum of these). Moreover, the valuation operators will depend on the particular firm; different firms will have different valuation operators, and an interesting question is how these combine.

In the next Section, we shall formulate the analogues of the axioms of Theorem 1 for *concave* valuation operators, and deduce some of their consequences. There are substantial differences; concave valuation operators have to be defined over cash balances, because without linearity we cannot build the price of a cash balance from the prices of its component parts. Nevertheless, the dynamic consistency axiom turns out to be the heart of the matter. We shall characterise families of valuation operators which satisfy the given axioms; it turns out that such families (and their duals) possess simple and appealing recursive structure.

We shall also study the question of how a firm may decide to divide up a risky cash balance process between its subsidiaries, each of which is subject to the regulatory constraints implicit in their individual valuation operators. We find that there is an optimal way to do this risk transfer, in terms of an inf-convolution (as in, for example, the study of Barrieu and El Karoui [4].) Moreover, the optimal risk transfer generates a family of valuation operators for the firm as a whole, and this family of valuation operators satisfies *the same axioms as the individual components*. We shall also see that if the firm decides at time 0 how it is going to divide up the cash balance between its subsidiaries, then at any later time, whatever has happened in the meantime, the original risk transfer chosen is still optimal. There is therefore a time-consistency in how the firm should transfer risk among its subsidiaries.

Another question we answer concerns what happens if a firm facing a risky cash balance process is allowed to take offsetting positions in a financial market. We find that there is an optimal offsetting position to be taken, which is time consistent, and the induced valuation operators for the firm once again satisfy the axioms.

2 Valuation operators.

Working in a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0\leq t\leq T}, P)$, we let BV denote the space of adapted processes of bounded variation with R-paths[8]. We think of $K \in BV$ as a *cash balance process*, with K_t being interpreted as the total amount of cash accumulated by time t. The process K need not of course be increasing. The upper end T of the time interval considered is a finite constant; there is no real difficulty in letting the time

[8]This is the terminology of Rogers & Williams [29] for paths that are right continuous with left limits everywhere.

set be $[0,\infty)$, but we choose not to do this here in view of our concentration later on examples where Ω is finite.

We propose to introduce some natural axioms to be satisfied by a family [9]

$$\{\pi_\tau : BV \to L^\infty(\mathcal{F}_\tau) \bigm| \tau \in \mathcal{O}\}$$

of *valuation operators*, or *valuations* for short. We interpret $-\pi_\tau(K)$ as the amount of capital required by law at time τ to allow a firm to accept the cash balance process K. The requirement could be different for firms in different countries, or for an investment bank and a hedge fund, for example. Not surprisingly, we shall suppose that

$$\pi_\tau(0) = 0 \quad \forall \tau. \tag{4}$$

The axioms we require of the family of valuations are the following.

(C) π_τ is concave for all τ;

(L) $\pi_\tau(I_A I_{[\tau,T]} K) = I_A \pi_\tau(K)$ for all τ, K, for all $A \in \mathcal{F}_\tau$;

(CL) if τ, τ' are two stopping times, and $A \in \mathcal{F}_\tau \cap \mathcal{F}_{\tau'}$, with $A \subseteq \{\tau = \tau'\}$, then for every K

$$\pi_\tau(K) = \pi_{\tau'}(K) \quad \text{on } A$$

(M) if $K_t \geq K'_t$ for all t, then $\pi_\tau(K) \geq \pi_\tau(K')$ for all τ;

(DC) for stopping times $\tau \leq \sigma$,

$$\pi_\tau(K) = \pi_\tau(K I_{[\tau,\sigma)} + \pi_\sigma(K) I_{[\sigma,\infty)});$$

(TI) if for some $a \in L^\infty(\mathcal{F}_\tau)$ we have $K_t = K'_t + a$ for all $t \geq \tau$, then $\pi_\tau(K) = a + \pi_\tau(K')$.

Remarks. Axiom (C) is a natural property for capital adequacy requirements for risky cash balances; see [2], for example.

Axiom (L) (for *local*) says two things. Firstly, if you have reached time τ, then all that matters for valuation is how much cash has *currently* been accumulated, and what is to come; the exact timing of the earlier payments does not influence the valuation[10]. Secondly, Axiom (L) expresses the following natural fact: at time τ, if event A has not happened then the cash balance $I_A I[0,\tau)K$ is clearly worthless, and if the event has happened, then the cash balance $I_A I[0,\tau)K$ will be worth the same as K. Axiom (CL)

[9] As usual, $\mathcal{O}$ denotes the optional σ-field on $[0,T]\times\Omega$, and by extension the statement $\tau \in \mathcal{O}$ for a random time τ means that $I_{[\tau,T]}$ is an optional process, equivalently, that τ is a stopping time - see, for example, [29] for more background on the general theory of processes.

[10] Note that Axiom (L) does *not* say that you value the cash balance after τ the same as the whole of the original cash balance K; the cash balance $I_{[\tau,T]}K$ pays nothing up til time τ, then a lump sum of K_τ.

(for *consistent localisation*) says that the localisations of π_τ and $\pi_{\tau'}$ agree where $\tau = \tau'$, again a natural condition.

Axiom (M) (for *monotonicity*) says that a larger capital reserve is required to short a larger cash balance, but it says more than just this. In particular, if $K_t \geq K'_t$ for all $0 \leq t \leq T$, with $K_T = K'_T$, then the two cash balances K and K' both deliver exactly the same in total, but K is considered less risky than K' because it delivers the cash *sooner*. An earlier version of this work used an axiom which expressed indifference between cash balances that delivered the same total amount of cash; though this axiom was entirely workable, the effect of it was that the valuation operators were essentially defined on cash balances which were all delivered at time T, and the valuation operators themselves served only to 'interpolate' prices in some sense. The interpretation of (M) is not that earlier payments are preferred to later payments because of the *interest* that will accrue; indeed, we think of all payments as being discounted back to time-0 values (or equivalently that the interest rate is zero). Even under these assumptions, according to (M) earlier payments are better than later ones - as in reality they are! This embodies the essence of cashflow problems, where a firm may be in difficulties not because it does not have sufficient money owed to it, but because that money has not yet come in.

Axiom (DC) (for *dynamic consistency*) has a simple and natural interpretation. It says that we must set aside as much for the cash balance K, as for the cash balance which gives us K up to time σ, and at time σ requires us to hand in the accumulated cash balance K_σ in return for the amount of cash that we would allow us to accept the entire cash balance K. This latter cash balance would clearly allow us to accept the original cash balance K.

The final axiom (TI) (for *translation invariance*) is again entirely natural.

In the next Section, we shall explore the consequences of these axioms only in the simplest possible setting, where Ω is *finite*. This means in particular that we can take the time set to be finite, and the entire filtered probability space to be represented by a tree. This (restrictive) assumption allows us to ignore all technicalities, and quickly uncover the essential structure implied by the axioms. We leave technicalities for technicians, and remark only that in any real-world application we would be forced to use a numerical approach, in which case we would have to be working with a finite sample-space.

3 Valuations on finite trees.

Henceforth, we work with a finite sample space Ω, and a finite time set $\{0, 1, \ldots, T\}$. The σ-field $\mathcal{F}$ on Ω is of course the σ-field of all subsets, and the filtration is represented by a tree[11] with vertex set $\mathcal{T}$. The root of the tree will be denoted by 0, and from any vertex $y \in \mathcal{T}$ there is a unique path to 0; we shall say that y is a descendant of x (written $x \preceq y$) if x lies on the path from y to 0. If $x \in \mathcal{T}$, we shall write $x - 1$ for the immediate ancestor of x, $x + 1$ for the set of immediate descendants of x, and $x+$ for the set of all descendants of x, including x itself. Note that Ω can be identified with the

[11]The tree does not of course have to be binomial, or regular.

set of endpoints of $\mathcal{T}$. For any $x \in \mathcal{T}$ we shall denote by $t(x)$ the *time* of x, which is the depth of x in the tree. Thus $t(0) = 0$, and $t(x) = T$ for any of the terminal nodes x. Notice also that a stopping time τ can be identified with a subset[12] $[\![\tau]\!]$ of $\mathcal{T}$ with the property that for any terminal node ω of the tree the unique path from ω to 0 intersects $[\![\tau]\!]$ in exactly one place.

In this setting, a *cash balance* is simply a map $K : \mathcal{T} \to \mathbb{R}$. We interpret K_x as the cumulative amount of the cash balance at vertex x in the tree. We shall also suppose throughout that interest rates are zero, or equivalently that cash balances have all been discounted back to time-0 values; this assumption is insubstantial, and leaves us clear to focus on what is important here.

In view of axiom (C), the valuations π_τ are just concave functions defined on some finite-dimensional Euclidean space, and so can be studied through their convex dual functions

$$\tilde{\pi}_\tau(\lambda) \equiv \sup_K \{\pi_\tau(K) - \lambda \cdot K\}.$$

For simplicity of exposition, we shall make the assumption

ASSUMPTION A: *For every τ, the valuation π_τ is concave, strictly increasing, upper semicontinuous, and C^2 in the relative interior of its domain of finiteness.*

By duality, the original functions π_τ can be expressed as

$$\pi_\tau(K) = \inf_\lambda \{\lambda \cdot K + \tilde{\pi}_\tau(\lambda)\}; \tag{5}$$

compare with the equation (2) above, as in Föllmer & Schied, Frittelli & Gianin. That equation is at one level simply the general statement (5) of duality, but with a bit more; in (2) the infimum is taken over a family of probability measures, and in (5) the infimum is unrestricted. We shall later see that the axioms used here do in fact imply that λ must be a probability on $x+$.

3.1 Decomposition.

The dynamic consistency axiom (DC) and localisation axioms (L), (CL) allow us to decompose the valuations in a simple way. To see this, notice firstly that the family (π_τ) of valuations is determined once the smaller family $\{\pi_x : x \in \mathcal{T}\}$ is known, where for $x \in \mathcal{T}$ the operator π_x is defined to be

$$\pi_x = \pi_{\tau_x}, \tag{6}$$

where τ_x is the stopping time

$$\begin{aligned} \tau_x(\omega) &= t(x) \quad \text{if } x \prec \omega; \\ &= T \quad \text{otherwise.} \end{aligned} \tag{7}$$

[12]The *graph* of τ - see [29].

Here, of course, we are identifying Ω with the set of terminal nodes of $\mathcal{T}$. Once we know the operators $\{\pi_x : x \in \mathcal{T}\}$, Axioms (L) and (CL) allow us to put together any of the $\pi_\mathcal{T}$.

However, the π_x can themselves be assembled from the family $\{\pi_{x,x+1} : x \in \mathcal{T}\}$ of *one-step valuation operators*, defined in the following way. If x is a terminal node, then the argument of $\pi_{x,x+1}$ is a cash balance k defined at x, and $\pi_{x,x+1}(k) = \pi_x(k)$. For all other x, given a cash balance k defined on $x \cup x+1$, we extend this to a cash balance $\bar{k}$ defined on all of $\mathcal{T}$ by

$$\begin{aligned} \bar{k}_z &= k_x \quad \text{if } z = x; \\ &= k_y \quad \text{if } y \preceq z \text{ for some } y \in x+1; \\ &= 0 \quad \text{otherwise.} \end{aligned}$$

We may then define

$$\pi_{x,x+1}(k) = \pi_{\tau_x}(\bar{k}). \tag{8}$$

Of course, the point of this decomposition is really the *converse*: we wish to build the (complicated) family $(\pi_\tau)_{\tau\in\mathcal{O}}$ from the simpler family $(\pi_{x,x+1})_{x\in\mathcal{T}}$ of one-step valuation operators. It is clear that if we derive $(\pi_{x,x+1})_{x\in\mathcal{T}}$ from a family $(\pi_\tau)_{\tau\in\mathcal{O}}$ satisfying the axioms given in Section 2, then the family of one-step valuations must have the following properties:

(c) $\pi_{x,x+1}$ is concave;

(m) if $k_z \geq k'_z$ for all $z \in x \cup x+1$ then $\pi_{x,x+1}(k) \geq \pi_{x,x+1}(k')$;

(ti) if $k_z = k'_z + a$ for all $z \in x \cup x+1$, then $\pi_{x,x+1}(k) = \pi_{x,x+1}(k') + a$ provided both are finite.

What we now argue is that given a family $(\pi_{x,x+1})_{x\in\mathcal{T}}$ of one-step valuation operators satisfying *(c), (m), (ti)* we can build a family $(\pi_\tau)_{\tau\in\mathcal{O}}$ of valuation operators satisfying the axioms of Section 2.

The essence of the construction is to get the $(\pi_x)_{x\in\mathcal{T}}$, for then if we have a stopping time τ we define

$$\pi_\tau(K) = \pi_z(K) \quad \text{at } z \in [\![\tau]\!].$$

To get the $(\pi_x)_{x\in\mathcal{T}}$, we proceed by backward induction, assuming that we have constructed π_x for all x such that $t(x) \geq n$. The induction starts, because if x is a terminal node we have $\pi_x(k) = \pi_{x,x+1}(k)$, and if $t(x) = n-1$ we may define

$$\pi_x(K) = \pi_{x,x+1}(k),$$

where k is the cash balance defined by

$$\begin{aligned} k_z &= K_x \quad \text{if } z = x; \\ &= \pi_z(K) \quad \text{if } z \in x+1. \end{aligned}$$

There is no problem with this, as the definition of π_x requires only the one-step operator $\pi_{x,x+1}$ and the operators $(\pi_z)_{z\in x+1}$ which are already known (by the inductive hypothesis).

If we vary the notation for $\pi_{x,x+1}(k) \equiv \pi_{x,x+1}(k_x, k_{x+1})$ so as to make the dependence on the cash balances at node x and nodes $x+1$ explicit, then the recursive construction of the π_x takes the clean form

$$\pi_x(K) = \pi_{x,x+1}(K_x, \pi_{x+1}(K)). \tag{9}$$

Notice the formal similarity to the notion of *recursive utility* - see Epstein & Zin [15], Duffie & Epstein [13], Skiadas [31]. This similarity is only formal; in the theory of recursive utility, there is a running consumption process which does not enter into our present discussion. This is an interesting extension of the axiomatic approach which we hope to return to at a later date.

It remains to see that the operators $(\pi_\tau)_{\tau\in\mathcal{O}}$ defined by (9) satisfy the axioms given in Section 2.

Property (C) follows from the concavity property *(c)* by backward induction. Properties (L) and (CL) are immediate from the construction. Property (M) follows from *(m)*, again by backward induction. Property (DC) requires a little more thought (and use of the property *(ti)*), but again follows from the construction. Finally, property (TI) is a consequence of *(ti)*.

3.2 Duality.

We have just seen that the axioms permit us to decompose the valuation operators into simpler pieces, but what is the corresponding result for the dual valuation operators $\tilde{\pi}_x$? What are the characteristic properties?

To understand the structure of the dual, firstly note that the dual valuation operator

$$\tilde{\pi}_x(\lambda) \equiv \sup_K \{\pi_x(K) - \lambda \cdot K\} \tag{10}$$

is not always going to be finite. Indeed, because of (L) and (CL), $\tilde{\pi}_x(\lambda)$ will be infinite if $\lambda_y \neq 0$ for some $y \notin x+$. Moreover, because of (M) the dual operator will be infinite if $\lambda_y < 0$ for some y. Finally, by considering cash balances K that are constant on $x+$ and using axiom (TI), we see that for finiteness of $\tilde{\pi}_x(\lambda)$ it is necessary that λ be a *probability* on $x+$: $\sum_{y\in x+} \lambda_y = 1$.

One further property can be deduced: $\inf_\lambda \tilde{\pi}_x(\lambda) = \pi_x(0) = 0$, using the duality relation and (4). Thus the dual valuation operators $(\tilde{\pi}_x)_{x\in\mathcal{T}}$ must satisfy the conditions

(D1) $\tilde{\pi}_x$ is convex;

(D2) $\tilde{\pi}_x(\lambda)$ is only finite if λ is a probability on $x+$;

(D3) $\inf_\lambda \tilde{\pi}_x(\lambda) = 0$.

The recursion for the dual valuations will follow from the recursive form (9) of the primal valuations. To make this explicit, we need to define the convex duals $\tilde{\pi}_{x,x+1}$ of the one-step valuation operators by the usual definition

$$\tilde{\pi}_{x,x+1}(\theta, \psi) = \sup_k \{ \, \pi_{x,x+1}(k_x, k_{x+1}) - \theta k_x - \psi \cdot k_{x+1} \}.$$

The analogues of (D1)-(D3) for the dual one-step valuations will be

(d1) $\tilde{\pi}_{x,x+1}$ is convex;

(d2) $\tilde{\pi}_{x,x+1}(\lambda)$ is only finite if λ is a probability on $x \cup x+1$;

(d3) $\inf_\lambda \tilde{\pi}_{x,x+1}(\lambda) = 0$.

It is easy to see that conditions (c), (m) and (ti) on the one-step valuation operators are equivalent to conditions (d1)-(d3) on their duals.

What then is the dual analogue of the primal recursion (9)? The answer is provided by the following result.

Theorem 2 *For all $x \in \mathcal{T}$ and λ a probability on $x+$, we have*

$$\tilde{\pi}_x(\lambda) = \tilde{\pi}_{x,x+1}(\lambda_x, \bar{\lambda}_{x+1}) + \sum_{z \in x+1} \bar{\lambda}_z \tilde{\pi}_z \left(\frac{\lambda_{\succeq z}}{\bar{\lambda}_z} \right) \tag{11}$$

where $\lambda_{\succeq z}$ denotes the restriction of λ to the set $\{y : y \succeq z\}$, and $\bar{\lambda}_z \equiv \sum_{y \succeq z} \lambda_y$.

Remark: Observe that the function

$$(\bar{\lambda}_z, \lambda_{\succeq z}) \mapsto \bar{\lambda}_z \tilde{\pi}_z \left(\frac{\lambda_{\succeq z}}{\bar{\lambda}_z} \right) = \sup \{ \bar{\lambda}_z \pi_z(K) - \lambda_{\succeq z} \cdot K \}$$

is convex.

Proof. Using (9) and (5), we have

$$\begin{aligned}
\pi_x(K) &= \inf_\lambda \{ \tilde{\pi}_x(\lambda) + \lambda \cdot K \} \\
&= \pi_x(K_x, \pi_{x+1}(K)) \\
&= \inf_{\lambda, \alpha} \{ \, \tilde{\pi}_{x,x+1}(\lambda_x, \alpha) + \lambda_x K_x + \alpha \cdot \pi_{x+1}(K) \, \} \\
&= \inf_{\lambda, \alpha, \psi} \{ \, \tilde{\pi}_{x,x+1}(\lambda_x, \alpha) + \lambda_x K_x + \alpha \cdot (\tilde{\pi}_{x+1}(\psi) + \psi \cdot K_{[x+1,T]}) \, \} \\
&= \inf_{\lambda, \alpha, \psi} \Big[\, \tilde{\pi}_{x,x+1}(\lambda_x, \alpha) + \lambda_x K_x + \sum_{z \in x+1} \alpha_z \{ \tilde{\pi}_z(\psi) + \psi_{\succeq z} \cdot K_{[z,T]} \} \, \Big] \\
&= \inf_{\lambda, \alpha} \Big[\, \tilde{\pi}_{x,x+1}(\lambda_x, \alpha) + \lambda_x K_x + \sum_{z \in x+1} \lambda_{\succeq z} \cdot K_{[z,T]} + \sum_{z \in x+1} \alpha_z \tilde{\pi}_z(\lambda_{\succeq z} / \alpha_z) \, \Big].
\end{aligned}$$

However, the only way that the terms inside the final infimum can be finite is if the arguments of the dual operators $\tilde{\pi}_{x,x+1}$ and $\tilde{\pi}_z$ are probabilities; and this only happens if $\alpha_z = \bar{\lambda}_z$ for all $z \in x+1$. The conclusion is that

$$\pi_x(K) = \inf_{\lambda} \Big[\ \tilde{\pi}_{x,x+1}(\lambda_x, \bar{\lambda}_{x+1}) + \sum_{z \in x+1} \bar{\lambda}_z \tilde{\pi}_z \Big(\frac{\lambda_{\succeq z}}{\bar{\lambda}_z} \Big) + \lambda \cdot K\ \Big]$$

and the result is proved. □

Notice that if we are given operators $(\tilde{\pi}_{x,x+1})_{x \in \mathcal{T}}$ satisfying (d1)-(d3), together with the condition that $\tilde{\pi}_{x,x+1} \equiv 0$ for any terminal node x, then the corresponding one-step valuations $(\pi_{x,x+1})_{x \in \mathcal{T}}$ satisfy (c), (m), (ti), and $\pi_{x,x+1}(k) = k_x$ for any terminal node x. By dualising (11) we quickly arrive at (9). Thus we may just as well construct a family of valuations operators $\pi_{\mathcal{T}}$ satisfying the axioms by starting from a family $(\tilde{\pi}_{x,x+1})_{x \in \mathcal{T}}$ of dual one-step valuations operators satisfying (d1)-(d3).

4 Examples

Let us consider some examples which can be analysed fairly completely in the tree setting.

Example 1: relative entropy. Suppose given some strictly positive probability distribution $(p_y)_{y \in \mathcal{T}}$ on $\mathcal{T}$. For any $x \in \mathcal{T}$ we define the dual valuation $\tilde{\pi}_x$ evaluated at some probability λ on $x+$ to be

$$\tilde{\pi}_x(\lambda) = \frac{1}{\gamma} \sum_{y \succeq x} \lambda_y \log(\lambda_y \bar{p}_x / p_y) \equiv h(\ \lambda_{\succeq x} \mid p_{\succeq x}/\bar{p}_x\ ), \tag{12}$$

where $\gamma > 0$ is some positive parameter, and as before $\bar{p}_x = \sum_{y \succeq x} p_y$. For other arguments, $\tilde{\pi}_x$ is infinite. It is well known that the function $\tilde{\pi}_x$ is convex, and its concave dual function is easily calculated to be

$$\pi_x(K) = -\frac{1}{\gamma} \log \Big[\sum_{y \succeq x} \frac{p_y}{\bar{p}_x} e^{-\gamma K_y} \Big], \tag{13}$$

equivalently,

$$e^{-\gamma \pi_x(K)} = \sum_{y \succeq x} \frac{p_y}{\bar{p}_x} e^{-\gamma K_y}. \tag{14}$$

It is now easy to check the axioms (C), (TP), (DC), (CE) of a family of valuation operators, and the axioms (L) and (CL) will hold by construction when we assemble the (π_x).

Remark. From (14) we might conjecture that similar examples could be constructed by the recipe

$$U(\pi_x(K)) = \sum_{y \succeq x} \frac{p_y}{\bar{p}_x} U(K_y)$$

for some other utility U. However, it is not clear that axioms (TI) and (C) will be satisfied in general, and indeed within some quite natural class the relative entropy example is the only example.

Example 2. This example is really a family of examples, built from the simple observation that if we have some collection $(\pi^\theta_{x,x+1})_{x\in\mathcal{T},\theta\in\Theta}$ of one-step valuation operators, such that for each θ the family $(\pi^\theta_{x,x+1})_{x\in\mathcal{T}}$ satisfies the axioms (c), (tp) and (ce), then the one-step valuation operators defined by

$$\pi_x(k) \equiv \inf_\theta \pi^\theta_{x,x+1}(k) \tag{15}$$

again satisfy (c), (tp) and (ce).

One simple example of this form could be constructed as follows. Suppose that for each $x \in \mathcal{T}$ we have some probability distribution $\alpha(x)$ on the immediate descendents $x+1$, and now we define

$$\pi_{x,x+1}(k) = \min\{k_x, \sum_{y\in x+1} \alpha(x)_y k_y\}.$$

The recursion (9) is now just the Bellman equation of dynamic programming, and the value of $\pi_x(K)$ is the 'worst stopping' value of the Markov decision process, where $\alpha(x)$ gives the distribution of moves down to $x+1$ from x if it is decided not to stop at x. It is easy to extend this example to the situation where a finite collection of possible distributions $\alpha^i(x)$ is considered at each vertex x, and the valuation operator gives the 'worst worst stopping' value!

Several of the examples of [6] are of this form, and we make no further remark on them. However, one feature that is noteworthy is the following. If we make the one-step valuation operators as infima of some sequence of linear functionals,

$$\pi_{x,x+1}(k) = \min_j \alpha^j \cdot k$$

then the valuations constructed are coherent, and the dual one-step valuation operators $\tilde{\pi}_x$ are

$$\begin{aligned}\tilde{\pi}_{x,x+1}(\lambda) &= 0 \quad \text{if } \lambda \in \mathrm{co}(\{\alpha^j\}) \\ &= \infty \quad \text{otherwise,}\end{aligned}$$

where $\mathrm{co}(A)$ denotes the convex hull of the set A. Looking at the recursive form (11) of the dual valuation operators, we see that these too take only the values 0 and ∞. The set of λ for which $\tilde{\pi}_0(\lambda)$ is finite is a *mutiplicatively stable* set.

Example 3. Families of one-step valuations $\pi_{x,x+1}$ can be constructed via the notion of a utility-indifference price for a single-period problem. In more detail, given some probability distribution p_y over $x \cup x+1$, we define $\pi_{x,x+1}(k)$ to be that value b such that

$$U(x_0) = \sum_{y\in x\cup x+1} p_y U(x_0 + k_y - b) \tag{16}$$

where U is some strictly increasing utility function, and x_0 is some reference wealth level. The properties (m) and (ti) are immediate, and (c) is a simple deduction.

It is unfortunately the case that there are few examples where the utility-indifference price for a single-period problem can be computed in closed form, and the dual valuation is similarly elusive. Some progress can be made however. Dropping the subscripts, the calculation of the dual valuation requires us to find

$$\tilde{\pi}(\lambda) = \sup_k \{\pi(k) - \lambda \cdot k\}$$

and the optimisation here can be considered as the optimisation

$$\sup_{b,k} \; b - \lambda \cdot k \tag{17}$$

$$\text{subject to} \quad U(x_0) = \sum p_y U(x_0 + k_y - b). \tag{18}$$

The Lagrangian form of the problem

$$\sup b, k\big[b - \lambda \cdot k + \theta(\sum p_y U(x_0 + k_y - b) - U(x_0))\big]$$

leads to the first-order conditions

$$\begin{aligned} 1 &= \theta \sum p_y U'(x_0 + k_y - b), \\ \lambda_y &= \theta p_y U'(x_0 + k_y - b), \end{aligned}$$

so (with $I \equiv (U')^{-1}$) we get

$$x_0 + k_y - b = I\big(\frac{\lambda_y}{\theta p_y}\big),$$

from which we see that θ is determined via

$$\sum p_y U\left(I\big(\frac{\lambda_y}{\theta p_y}\big) \right) = U(x_0).$$

The final expression

$$\tilde{\pi}(\lambda) = x_0 - \sum \lambda_y I\left(\frac{\lambda_y}{\theta p_y}\right)$$

simplifies in the case of CRRA $U(x) = x^{1-R}/(1-R)$ to

$$\tilde{\pi}(\lambda) = x_0 - x_0^R \left(\sum p_y^{1/R} \lambda_y^{1-1/R} \right)^{R/(R-1)}.$$

5 Spreading and evolution of risk.

Let us consider the situation of a firm which consists of J subsidiaries, possibly in different countries, or subject to different regulatory controls. We let the valuation operators $(\pi^j_\tau)_{\tau\in\mathcal{O}}$ determine the regulatory requirements of subsidiary j, $j = 1, \ldots, J$. If (at τ) subsidiary i wishes to accept the cash balance process K, then regulation requires that subsidiary to reserve $-\pi^i_\tau(K)$. However, subsidiary i could approach another subsidiary j and get them to take from i the cash-balance process K^j in return for the regulatory capital $-\pi^j_\tau(K^j)$. Subsidiary i is free to enter into such agreements with all the other subsidiaries, and will do so in such a way as to minimise the regulatory capital required. Taking into account the possibilities of risk transfer, subsidiary i will need to reserve $-\Pi^i_\tau(K)$ instead of $\pi^i_\tau(K)$, where

$$\begin{aligned}
\Pi^i_\tau(K) &= \sup\{\pi^i_\tau(K - \sum_{j\neq i} K^j + \sum_{j\neq i} \pi^j_\tau(K^j) I_{[x,T]})\} \\
&= \sup\{\pi^i_\tau(K - \sum_{j\neq i} K^j) + \sum_{j\neq i} \pi^j_\tau(K^j)\} \\
&= \sup\{\sum_j \pi^j_\tau(K^j) : \sum_j K^j = K\}. \qquad (19)
\end{aligned}$$

Notice that this is *independent of the choice of subsidiary*, so we write simply Π_τ for Π^i_τ. Moreover, we may have that $\Pi_\tau(0) > 0$, so we define

$$^0\Pi_\tau(K) = \Pi_\tau(K) - \Pi_\tau(0), \qquad (20)$$

so as to have the property (4) for the operators $^0\Pi_\tau$. The quantity $\Pi_x(0)$ can be interpreted as the *value of risk-sharing* at vertex x. We call the family $(^0\Pi_\tau)_{\tau\in\mathcal{O}}$ of valuations the *risk-sharing* valuation operators, though it is not clear as yet that we may refer to them as such, since we do not know that they satisfy the axioms for a family of valuation operators. That is the task of the following result.

Theorem 3 *The risk-sharing valuations $(^0\Pi_\tau)_{\tau\in\mathcal{O}}$ satisfy the axioms (C), (L), (CL), (M), (DC), and (TI) of the component valuations $(\pi^j_\tau)_{\tau\in\mathcal{O}}$, $j = 1, \ldots, J$.*

Proof. Properties (C), (L), (CL), (M), and (TI) are straightforward to verify; only the property (DC) is not immediately obvious. To establish this, we have on the one hand

$$\begin{aligned}
\Pi_x(K) &= \sup\{\sum_j \pi^j_x(K^j) : \sum_j K^j = K\} \\
&= \sup\{\sum_j \pi^j_x(K^j I_{[x,\tau)} + \pi^j_\tau(K^j) I_{[\tau,T]}) : \sum_j K^j = K\} \\
&= \sup\{\sum_j \pi^j_x(K^j I_{[x,\tau)} + a_j I_{[\tau,T]}) : \sum_j K^j I_{[x,\tau)} = K I_{[x,\tau)}, \\
&\qquad \sum_j a_j \leq \Pi_\tau(K)\} \qquad (21)
\end{aligned}$$

and on the other hand we have

$$\begin{aligned}
\Pi_x(K_{[x,\tau)} + {}^0\Pi_\tau(K)I_{[\tau,T]}) &= \sup\{\sum_j \pi_x^j(K^j) : \sum_j K^j = K_{[x,\tau)} + {}^0\Pi_\tau(K)I_{[\tau,T]}\} \\
&= \sup\{\sum_j \pi_x^j(K^j I_{[x,\tau)} + \pi_\tau^j(K^j)I_{[\tau,T]}) : \\
&\qquad\qquad \sum_j K^j = K_{[x,\tau)} + {}^0\Pi_\tau(K)I_{[\tau,T]}\} \\
&= \sup\{\sum_j \pi_x^j(K^j I_{[x,\tau)} + a_j I_{[\tau,T]}) : \sum_j K^j I_{[x,\tau)} = K I_{[x,\tau)}, \\
&\qquad\qquad \sum_j a_j \leq \Pi_\tau({}^0\Pi_\tau(K)I_{[\tau,T]})\} \quad (22)
\end{aligned}$$

But $\Pi_\tau({}^0\Pi_\tau(K)I_{[\tau,T]}) = \Pi_\tau(0) + {}^0\Pi_\tau(K) = \Pi_\tau(K)$ and comparing (21) and (22) we see that $\Pi_x(K) = \Pi_x(K_{[x,\tau)} + {}^0\Pi_\tau(K)I_{[\tau,T]})$, equivalently, ${}^0\Pi_x(K) = {}^0\Pi_x(K_{[x,\tau)} + {}^0\Pi_\tau(K)I_{[\tau,T]})$, as required.

∎

There is a simple interpretation of risk-sharing in terms of the duals. Indeed, from (19) we have that

$$\begin{aligned}
\tilde{\Pi}_x(\lambda) &= \sup_{(K^j)}\{\sum_j \pi_x^j(K^j) - \lambda \cdot \sum_j K^j\} \\
&= \sum_j \tilde{\pi}_x^j(\lambda),
\end{aligned}$$

so the effect of risk-sharing is simply to *add the dual valuation operators.*

5.1 Optimal risk transfer in the relative entropy example.

If each of the J subsidiaries has valuations of the relative entropy form (recall (13)):

$$e^{-\gamma_j \pi_x^j(K)} = \sum_{y \succeq x} \frac{p_y^j}{\bar{p}_x^j} e^{-\gamma_j K_y}, \tag{23}$$

how do they combine under risk sharing? For the moment, let us fix a particular $x \in \mathcal{T}$ and consider how things work from that node. We shall write $\tilde{p}_y^j \equiv p_y^j/\bar{p}_x^j$ for brevity, and shall define

$$\Gamma \equiv \big(\sum_j \gamma_j^{-1}\big)^{-1}. \tag{24}$$

The dual valuation operators are given by (see (12))

$$\tilde{\pi}_x^j(\lambda) = \frac{1}{\gamma_j} \sum_{y \succeq x} \lambda_y \log(\lambda_y/\tilde{p}_y^j),$$

so the risk-sharing result gives us

$$
\begin{aligned}
\tilde{\Pi}_x(\lambda) &= \sum_j \tilde{\pi}_x^j(\lambda) \\
&= \sum_j \frac{1}{\gamma_j} \sum_{y \succeq x} \lambda_y \log(\lambda_y / \tilde{p}_y^j) \\
&= \frac{1}{\Gamma} \Big\{ \sum_{y \succeq x} \lambda_y \log \lambda_y - \sum_j \frac{\Gamma}{\gamma_j} \sum_{y \succeq x} \lambda_y \log \tilde{p}_y^j \Big\} \\
&= \frac{1}{\Gamma} \sum_{y \succeq x} \lambda_y \log(\lambda_y / P_y) - \frac{1}{\Gamma} \log \Big\{ \sum_{y \succeq x} \prod_i (p_y^i)^{\Gamma/\gamma_i} \Big\}
\end{aligned}
$$

where we define the probability P on $x+$ by

$$
P_y \equiv \frac{\prod_i (p_y^i)^{\Gamma/\gamma_i}}{\sum_{z \succeq x} \prod_i (p_y^i)^{\Gamma/\gamma_i}}. \tag{25}
$$

From this we see that

$$
\begin{aligned}
\Pi_x(0) &= -\frac{1}{\Gamma} \log \Big\{ \sum_{y \succeq x} \prod_i (p_y^i)^{\Gamma/\gamma_i} \Big\} \\
&= -\frac{1}{\Gamma} \log \Big\{ \sum_{y \succeq x} \exp\Big(\sum_j \frac{\Gamma}{\gamma_j} \log(\tilde{p}_y^j)\Big) \Big\} \\
&\geq 0,
\end{aligned}
$$

by Jensen's inequality, with equality if and only if all the agents have the same p_y^j. Thus we see that the aggregated dual valuations ${}^0\tilde{\Pi}_x$ have the *same* relative-entropy form as the individual dual valuations, with explicit expressions (24) for the combined coefficient of absolute risk aversion Γ and (25) for the combined distribution of the probability down the tree.

How does the risk sharing work out in this example? The maximisation (19) of $\sum_j \pi_x^\tau(K^j)$ can be computed, leading to the conclusion that

$$
\begin{aligned}
K_y^j &= \frac{\Gamma}{\gamma_j} K_y + \Big\{ \frac{1}{\gamma_j} \log \tilde{p}_y^j - \frac{\Gamma}{\gamma_j} \Big(\sum_i \frac{1}{\gamma_i} \log \tilde{p}_y^i \Big) \Big\}. & (26) \\
&= \frac{\Gamma}{\gamma_j} K_y + \frac{1}{\gamma_j} \log(\tilde{p}_y^j / P_y) + \frac{\Gamma}{\gamma_j} \Pi_x(0) & (27)
\end{aligned}
$$

This provides a nice interpretation of the way that the cash balance K gets shared. At each node y, the cash balance K_y at the node gets split proportionally between the subsidiaries ('linear risk sharing'), and there are a further two terms, one relating to the ratio of subsidiary j's probability of the node y and the aggregated probability P_y, and the other proportional to $\Pi_y(0)$.

5.2 Dynamic stability of the risk-sharing solution.

When computing the value $\Pi_0(0)$ of risk-sharing at time 0, the subsidiaries find themselves solving the optimisation problem

$$\sup\{\sum_j \pi_0^j(K^j) : \sum_j K^j = 0\}.$$

Casting the problem in Lagrangian form

$$\sup\{\sum_j [\,\pi_0^j(K^j) - p \cdot K^j\,],$$

it is easy to see that at an optimal solution we shall have that all subsidiaries' marginal valuations of cash balances will coincide:

$$\nabla \pi_0^j(K^j) = p. \tag{28}$$

Suppose that at time 0 they adopt the optimal cash balance processes K^j obtained in this way; as time passes, *will they still be satisfied with the K^j they first agreed to?* It would be disturbing if we reached some vertex x in the tree where the subsidiaries would wish to renegotiate the deals that they had committed to at time 0. However, it turns out that *this does not happen*: and it is the condition (DC) and the chain rule which guarantees this.

If x is some vertex in the tree, and we let $\tau = \tau_x$ (recall (7)), then using (DC) we have

$$\pi_0^j(K) = \pi_0^j(KI_{[0,\tau)} + \pi_\tau(K)I_{[\tau,T]})$$

and differentiating both sides with respect to K_y, where $y \succeq x$, gives us (by the chain rule)

$$p_y = \frac{\partial \pi_0^j}{\partial K_y}(K^j) = \frac{\partial \pi_0^j}{\partial K_x}(K^j I_{[0,\tau)} + \pi_\tau(K^j)I_{[\tau,T]})\frac{\partial \pi_x^j}{\partial K_y}(K).$$

Accordingly, in view of (28), we have for each j that there exists a constant b_j such that for all $y \succeq x$

$$\frac{\partial \pi_x^j}{\partial K_y}(K) = b_j p_y,$$

and so at vertex x the remaining allocations (cash balances) $K^j I_{[x,T]}$ still constitute a competitive equilibrium; there are no mutually beneficial trades available to the agents at vertex x.

Remarks. We could have discussed this dynamic stability in terms of competitive equilibria. Indeed, if we were to write $U^j(K)$ in place of $\pi_0^j(K)$, then the concave increasing functions U^j can serve as the utilities of different agents, defined over bundles of goods, where cash balances at different vertices are interpreted as different goods. We are now in the realm of finding an equilibrium allocation, and this is dealt with in any decent text on microeconomic theory; see, for example, [24]. However, although the mathematics is exactly that of finding an equilibrium in a pure exchange economy, such an analogy is economically impure; here, we have been interpreting the π_x^j as some sort of *price*, not as a utility. We allow (for example) the values $\pi_x^j(K)$ into the *arguments* of the functions (utilities?!) π_0^j. It seems to us that the link is tenuous, and we desist from pushing the analogy too far.

5.3 Spreading risk by access to a market.

Suppose a firm with valuation operators $(\pi_x)_{x\in\mathcal{T}}$ is allowed access to a market; how will it act, and how does its valuation of cash balances change? The discussion is similar to that of risk sharing among subsidiaries, but sufficiently different to require a separate treatment.

We represent access to the market in the following way. At each stopping time τ, the firm may change a given cash balance process K to $K+K'$ for any $K' \in G_\tau$, where G_τ denotes the gains-from-trade cash balance processes which could be achieved by trading in the market starting with zero wealth at time τ. Concerning the G_τ we shall assume[13] that

(c-m) each G_τ is convex;

(l-m) for each $x \in [\![\tau]\!]$,

$$G_{\tau_x} = \{KI_{[x,T]} : K \in G_\tau\};$$

(dc-m) for each $\tau \le \sigma \in \mathcal{O}$ if K^σ denotes[14] the cash balance process K stopped at σ, we have

$$G_\tau = \{K^\sigma + K' : K \in G_\tau, K' \in G_\sigma\}$$

Now the cash balance valuation given access to this market will be via

$$\Pi_x(K) \equiv \sup\{\pi_x(K+K') : K' \in G_x\}. \tag{29}$$

Once again, there is no guarantee that $\Pi_x(0) = 0$, but if we define

$${}^0\Pi_x(K) \equiv \Pi_x(K) - \Pi_x(0), \tag{30}$$

then the operators $({}^0\Pi_x)_{x\in\mathcal{T}}$ do have this property (4). As in the case of risk-sharing, the quantity $\Pi_x(0)$ is the value to the agent of being granted access to the market at vertex x.

Theorem 4 *The valuations $({}^0\Pi_\tau)_{\tau\in\mathcal{O}}$ satisfy the axioms (C), (L), (CL), (M), (DC), and (TI).*

[13] Recall the definition (7) of τ_x.

[14] Formally, $K_z^\sigma = K_y$ if $z \succeq y \in [\![\sigma]\!]$; $= K_z$ otherwise.

Proof. As before, all of the properties except for (DC) are obvious. To prove (DC), we use properties (DC) for the π_x and (dc-m) to develop

$$\begin{aligned}\Pi_x(K) &= \sup\{\pi_x(K+K') : K' \in G_x\} \\ &= \sup\{\pi_x((K+K')I_{[x,\sigma)} + \pi_\sigma(K+K')I_{[\sigma,T]}) : K' \in G_x\} \\ &= \sup\{\pi_x((K+K')I_{[x,\sigma)} + \{\Pi_\sigma(K) + K'_\sigma\}I_{[\sigma,T]}) : K' \in G_x\}. \end{aligned} \tag{31}$$

On the other hand,

$$\begin{aligned}\Pi_x(KI_{[x,\sigma)} + {}^0\Pi_\sigma(K)I_{[\sigma,T]}) &= \sup\{\pi_x(KI_{[x,\sigma)} + {}^0\Pi_\sigma(K)I_{[\sigma,T]} + K') : K' \in G_x\} \\ &= \sup\{\pi_x((K+K')I_{[x,\sigma)} + {}^0\Pi_\sigma(K)I_{[\sigma,T]} + K'I_{[\sigma,T]}) : K' \in G_x\} \\ &= \sup\{\pi_x(\ (K+K')I_{[x,\sigma)} + {}^0\Pi_\sigma(K)I_{[\sigma,T]} + \\ &\qquad + \pi_\sigma(K'I_{[\sigma,T]})I_{[\sigma,T]}\ ) : K' \in G_x\} \\ &= \sup\{\pi_x(\ (K+K')I_{[x,\sigma)} + {}^0\Pi_\sigma(K)I_{[\sigma,T]} + \\ &\qquad + (\Pi_\sigma(0) + K'_\sigma)I_{[\sigma,T]}\ ) : K' \in G_x\} \end{aligned} \tag{32}$$

since when we maximise $\pi_\sigma(K'I_{[\sigma,T]})$ over K' we get $\Pi_\sigma(0) + K'_\sigma$. Comparing (31) and (32) establishes property (DC) for the operators $({}^0\Pi_\tau)_{\tau\in\mathcal{T}}$. ∎

6 Conclusions.

This paper has approached the problem of convex risk measurement in a dynamic setting from a slightly unconventional starting point; instead of trying to work with acceptance sets, we begin with valuation operators satisfying certain axioms which seem to us to be natural. Our notion of preference does *not* reduce to a simple valuation of all the proceeds of the cashflow collected at the end, but genuinely accounts for the (obvious) fact that you would prefer to have \$1M today rather than the value of \$1M invested at riskless rate in five years from now.

In the simplest situation, where the sample-space is finite, we show how a family of pricing operators obeying our axioms can be decomposed into (and reconstructed from) a family of one-period pricing operators which are much easier to grasp. There is a corresponding decomposition of the dual pricing functions.

Allowing a firm to spread risk among a number of subsidiaries leads to risk-sharing solutions; the firm derives benefit from risk sharing, and, remarkably, the risk-sharing valuations which arise satisfy exactly the same set of axioms satisfied by the initial valuations.

We have seen also that the risk sharing that arises will be stable over time; if at time 0 the firm chooses how to spread risk among its subsidiaries, then no matter how the world evolves, at all later times it will continue to be satisfied with the cash balances that it originally selected.

We study also what happens when a firm is allowed access to a financial market. Assuming some natural properties of the market, the conclusions are similar to the risk-sharing problem; the firm derives a fixed benefit from being allow access to the market, but beyond that it values cash balance processes according to modified valuation operators which satisfy the same axioms.

A Appendix.

A.1 Proof of Theorem 1.

Let us consider the filtered probability space $(\Omega, (\mathcal{F}_t)_{t\geq 0}, P)$. For any $T > 0$, the map $A \mapsto \pi_{0T}(\mathbf{I}_A)$ defines a non-negative measure on the σ-field $\mathcal{F}_T$, from the linearity, positivity and continuity properties of our pricing operator.

This measure is moreover absolutely continuous with respect to $\mathbf{P}$ in view of (A2). Hence by the Radon-Nikodym theorem, there exists a non-negative $\mathcal{F}_T$-measurable random variable ζ_T such that

$$\pi_{0T}(Y) = \mathbf{E}[\zeta_T Y]$$

for all $Y \in L^\infty(\mathcal{F}_T)$.

Moreover, (A2) implies that $\mathbf{P}[\zeta_T > 0] = 1$.

We finally use the consistency condition (A3) as follows. Let $Y \in L^\infty(\mathcal{F}_T)$, then by definition, $\pi_{tT}(Y) \in L^\infty(\mathcal{F}_t)$. For any $X \in L^\infty(\mathcal{F}_t)$,

$$\begin{aligned}
\pi_{0t}(X\pi_{tT}(Y)) &= \mathbf{E}[X\zeta_t\pi_{tT}(Y)] \\
&= \pi_{0T}(XY) \\
&= \mathbf{E}[XY\zeta_T]
\end{aligned}$$

Since $X \in L^\infty(\mathcal{F}_t)$ is arbitrary, we deduce that

$$\pi_{tT}(Y) = \frac{1}{\zeta_t}\mathbf{E}_t[\zeta_T Y]$$

which shows that the pricing operators π_{st} are actually given by the risk-neutral pricing recipe (3) described in Theorem 1, with the state-price density process ζ.

The state-price density process is often thought of as the product of the discount factor $\exp\left(-\int_0^t r_s \mathrm{d}s\right)$ and the change-of-measure martingale.

References

[1] P. Artzner, F. Delbaen, J.M. Eber, and D. Heath. Thinking coherently. *Risk*, 10:68–71, 1997.

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

[3] P. Artzner, F. Delbaen, J.M. Eber, D. Heath, and H. Ku. Coherent multiperiod risk adjusted values and Bellman's principle. *Preprint*, 2004.

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

[5] P. Cheredito, F. Delbaen, and M. Kupper. Coherent and convex risk measures for bounded cadlag processes. *WP ETH, Zürich*, 2003.

[6] P. Cheredito, F. Delbaen, and M. Kupper. Dynamic monetary risk measures for bounded discrete-time processes. *WP ETH, Zürich*, 2004.

[7] A. S. Cherny. Pricing, optimality, and equilibrium based on coherent risk measures. *Preprint*, 2005.

[8] J. Cvitanič and I. Karatzas. On dynamic measures of risk. *Finance and Stochastics*, 3:451–482, 1999.

[9] F. Delbaen. Coherent risk measures. Lectures given at the Cattedra Galileiana. *Scuola Normale Superiore, Pisa*, 2000.

[10] F. Delbaen. Coherent risk measures on general probability spaces. *Advances in Finance and Stochastics, Springer, New York*, 2002.

[11] F. Delbaen. The structure of m-stable sets and in particular of the set of risk-neutral measures. *WP ETH, Zürich*, 2003.

[12] O. Deprez and H. Gerber. On convex principles of premium calculation. *Insurance: Mathematics and Economics*, 4:179–189, 1985.

[13] D. Duffie and L. Epstein. Asset pricing with stochastic differential utility. *The Review of Financial Studies*, 5:411–436, 1992.

[14] L. Epstein and M. Schneider. Recursive multiple priors. *Journal of Economic Theory*, 113:1–31, 2003.

[15] L. Epstein and S. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica*, 57:937–969, 1989.

[16] H. Föllmer and A. Schied. Convex measures of risk and trading constraints. *Finance and Stochastics*, 6:429–447, 2002.

[17] H. Föllmer and A. Schied. *Stochastic Finance.* de Gruyter, Berlin, 2002.

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

[19] M. Frittelli and E. Rosazza Gianin. *Dynamic convex risk measures.* Risk measures for the 21st century, Wiley, 2004.

[20] M. Frittelli and E. Rosazza Gianin. Law invariant convex risk measures. *Advances in Mathematical Economics*, 7:33–46, 2005.

[21] M. Frittelli and G. Scandolo. Risk measures and capital requirements for processes. *Technical Report 6, Department Matematica per le Decisioni, University of Firenze*, 2005.

[22] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

[23] Jaschke and Küchler. Coherent risk measures and good-deal bounds. *Finance and Stochastics*, 5:181–200, 2001.

[24] A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory.* OUP, 1995.

[25] Y. Nakano. Minimizing coherent risk measures of shortfall in discrete-time models with cone constraints. *Applied Mathematical Finance*, 10:163–181, 2003.

[26] S. Peng. Backward stochastic differential equations and related g-expectation. *Pitman Research Notes in Mathematics Series*, 364:141–159, 1997.

[27] S. Peng. Dynamically consistent nonlinear evaluations and expectations. *Preprint*, 2005.

[28] F. Riedel. Dynamic coherent risk measures. *Working Paper, Stanford University*, 2003.

[29] L. C. G. Rogers and D. Williams. *Diffusions, Markov Processes, and Martingales, Volume 2.* Cambridge University Press, Cambridge, 2000.

[30] L.C.G. Rogers. The origins of risk-neutral pricing and the Black-Scholes formula. *Risk Management and Analysis*, 2:81–94, 1998.

[31] C. Skiadas. Recursive utility and preferences for information. *Economic Theory*, 12:293–312, 1998.