Dynamic utility-based good deal bounds

Susanne Klöppel, Martin Schweizer

Summary: We introduce and study no-good-deal valuation bounds defined in terms of expected utility. A utility-based good deal is a payoff whose expected utility is too high in comparison to the utility of its price. Forbidding good deals induces, via duality, restrictions on pricing kernels and thereby gives tighter valuation bounds on payoffs than absence of arbitrage alone. Our approach extends earlier work by Černý (2003) in several directions: We give rigorous results for a general probability space instead of finite Ω; we systematically use duality results to provide a streamlined approach with simple arguments; we do all this rigorously for both static and dynamic situations; and we give a systematic comparison between local and global (conditional) pricing kernel restrictions for the temporally dynamic setting. For the dynamic case, we show in a Lévy framework that defining no-good-deal valuation measures by imposing local conditional restrictions on their instantaneous market prices of risk gives valuation bounds having very good dynamic properties as processes over time. We also show that global restrictions cannot yield such results in general.

1 Introduction

An *arbitrage opportunity*, in loose terms, is a tradable payoff producing a sure gain with no risk of loss. By the fundamental theorem of asset pricing (FTAP), (a precise form of) absence of arbitrage opportunities is equivalent to the existence of an equivalent local martingale measure (ELMM) for the traded assets in the given financial market. Being a preference-free concept, absence of arbitrage is a mild and universally accepted assumption, but its implications for pricing are often rather weak. In an incomplete market, the range of arbitrage-free prices for a payoff X is the full interval from inf $E_Q[X]$ to sup $E_Q[X]$, where Q runs through all ELMMs, and these two bounds are usually too far apart to give useful information.

Sharper bounds for payoff values can be obtained via preferences. These allow to quantify the attractiveness of payoffs in order to forbid those, called *good deals*, which are "too good to be true". So a good deal is a quantitative sharpening of the qualitative concept of an arbitrage opportunity. No-good-deal valuation bounds are then obtained by assuming absence of good deals in the market, and the bounds again take the form inf $E_Q[X]$ and sup $E_Q[X]$, but Q now only runs through a subset $\mathcal{N}$ of all ELMMs. These

no-good-deal measures are obtained by translating absence of good deals via duality into a restriction on pricing kernels $\frac{dQ}{dP}$.

The best-known starting point for good deal bounds is Cochrane and Saá-Requejo (2000); the simultaneously published paper by Bernardo and Ledoit (2000) is quoted less frequently. Cochrane and Saá-Requejo (2000) measure the attractiveness of a payoff by its Sharpe ratio and use the Hansen–Jagannathan inequality to find as corresponding restriction a bound on the variance of the pricing kernel. This is done in a static (i.e., one-period) setting and extended with mainly formal limit arguments to diffusions. The latter part has been streamlined and extended to models with jumps by Björk and Slinko (2006); see Section 3 for more details.

The same idea as in Cochrane and Saá-Requejo (2000) is developed in Bernardo and Ledoit (2000). Their criterion is the *gain-loss ratio* $E_R[X^+]\big/E_R[X^-]$ for a pricing measure R and for X from the set $\mathcal{C}(0, S)$ of all payoffs that can be superreplicated at zero initial cost from the primary assets S. The key result in Bernardo and Ledoit (2000) is the duality

$$\sup_{X\in\mathcal{C}(0,S)} \frac{E_R[X^+]}{E_R[X^-]} = \inf_{Q\in\{\text{ELMMs}\}} \frac{\operatorname{ess\,sup}\frac{dQ}{dR}}{\operatorname{ess\,inf}\frac{dQ}{dR}},$$

proved for finite Ω and used as in Cochrane and Saá-Requejo (2000) to obtain good deal valuation bounds from a restriction on Q. The two approaches are neatly summarised by Longarela (2001) who also points out their general structure: To obtain good deal bounds, one needs on the one hand criteria to quantify payoffs' performances and to measure pricing kernel variability, and on the other hand a duality result linking these two. A second contribution of Longarela (2001) is to propose a further criterion (on the Q-side) and to prove the corresponding duality; see also Bondarenko and Longarela (2004) for a slightly different presentation. Another overview of results and ideas for bounding pricing kernels appears in Ross (2005), but puts less emphasis on the application to good deal bounds.

Conceptually, an important step forward is then taken by Černý (2003), for a *finite* probability space Ω. After observing that the Sharpe ratio bounds can equivalently be deduced by working with quadratic utility, he proposes to quantify payoffs' attractiveness through a utility function U via a notion of *generalised Sharpe ratios*. In essence, this boils down to comparing the U-certainty equivalent $U^{-1}(E[U(X)])$ with the Q-price $E_Q[X]$ of X. Černý (2003) introduces an interval of *(utility-based) no-good-deal values* and gives explicit formulae for pricing kernel restrictions for a number of well-known utility functions. All this builds on, and significantly extends, many early ideas in Hodges (1998). But despite important conceptual contributions, we also see in Černý (2003) a number of drawbacks. Proofs are only given for finite Ω, and several crucially need this restriction. There are extensions to a dynamic setting with Itô processes, but they are obtained via heuristic or formal limit arguments and in particular inherently cannot deal with continuous-time models where S has jumps.

The present paper makes three contributions. We first introduce and represent utility-based no-good-deal valuation bounds in a static framework for a general probability space, with rigorous proofs. The key tool is a fundamental *duality* from Kramkov and Schachermayer (1999) between utility maximisation and optimal martingale measures. In particular, this brings up pricing kernel restrictions via duality in a clear, systematic and rigorous way.

Secondly, we present a new approach to a temporally *dynamic* setting and discuss how to choose there bounds on pricing kernels. The limit arguments in Cochrane and Saá-Requejo (2000) and Černý (2003) only cover diffusions and automatically but heuristically lead to pricing kernel bounds in the form of (conditional) *local and quadratic restrictions* on the instantaneous market price of risk. Björk and Slinko (2006) extend this rigorously and even include jumps from a marked point process, but still keep the framework of quadratic restrictions corresponding to conditional instantaneous Sharpe ratios. For our utility-based approach, we start instead in the setting of a Lévy filtration and study both global and local conditional restrictions. Using results from Delbaen (2006), we show that our dynamic good deal valuation bounds obtained from local restrictions have very natural and nice properties as processes over time, and that this fails in general for global instead of local restrictions on pricing kernels.

Our third contribution has two parts. We give sufficient conditions for the set of dynamic no-good-deal valuation measures to be nonempty, and we provide a detailed discussion of the connections between local and global restrictions in a Lévy setting.

The paper is structured as follows. Section 2 explains the basic ideas in the static (one-period) case, going from Sharpe ratios over general utility functions to specific examples of sets $\mathcal{N}$ of no-good-deal measures. Section 3 contains the dynamic treatment of no-good-deal bounds. Section 4 gives an overview of abstract good deals defined via risk measures, including a discussion on that slightly different strand of literature and linking it to the utility-based approach. Section 5 concludes with a summary and outlook. The Appendix collects results on Lévy processes used in Section 3.

2 The basic ideas in the static case

This section presents the main ideas in the simplest, static setting. After reviewing classical no-good-deal bounds from Sharpe ratios, we explain our new method which combines (von Neumann–Morgenstern) utility with a systematic use of duality results.

We start with some basic concepts. A *payoff* is a random variable X on a probability space $(\Omega, \mathcal{F}, P)$; it gives at maturity T in state $\omega \in \Omega$ the amount $X(\omega)$ in units of a fixed numeraire. A *pricing measure* is a probability measure $Q \approx P$, briefly $Q \in \mathcal{P}^{\approx}$. For initial capital $x \in \mathbb{R}$ and a pricing measure Q, the set of *Q-affordable* payoffs is

$$\mathcal{C}(x, Q) := \left\{ X \in L^0 \mid X^- \in L^\infty \text{ and } E_Q[X] \leq x \right\}.$$

Having X bounded from below is natural for payoffs. We call Q an L^2-*pricing measure* if its *pricing kernel* $Z_T := \frac{dQ}{dP}$ is in $L^2(P)$. Unspecified expectations and variances are under P.

Definition 2.1 *For $X \in L^2(P)$ and an L^2-pricing measure Q, the (*unconditional, global*) Sharpe ratio (over $[0, T]$) is the return over risk quotient*

$$SR(X, Q) := \frac{E[X] - E_Q[X]}{\sqrt{\text{Var}[X]}}.$$

If X is constant, or if $\text{Var}[X] = \infty$, we set $SR(X, Q) = 0$.

The Sharpe ratio is a popular performance measure widely used in the literature. Since empirically observed Sharpe ratios tend to be rather low, we follow Cochrane and Saá-Requejo (2000) and call a payoff $X \in L^2(P)$ a *classical (Q-)good deal of level δ* if $SR(X, Q) \geq \delta$. Observe that $SR(X, Q)$ involves the Q-price $E_Q[X]$ of X in its definition.

The idea in *no-good-deal valuation* is to use only pricing measures that do not turn any payoff into a "too good" deal. To make this quantitative, one classically writes the numerator of $SR(X, Q)$ as $E[(X - E[X])(1 - Z_T)]$ and applies Cauchy–Schwarz to obtain

$$SR(X, Q) \leq \sqrt{\mathrm{Var}[Z_T]} \text{ for every payoff } X \in L^2(P), \qquad (2.1)$$

the *Hansen–Jagannathan inequality*. With a little extra work, one can sharpen (2.1) to

$$\sup_{\substack{X \in \mathcal{C}(0,Q), \\ E[X] < \infty}} SR(X, Q) = \sup_{\substack{X \in \mathcal{C}(x,Q), \\ E[X] < \infty}} SR(X, Q) = \sqrt{\mathrm{Var}[Z_T]} \qquad (2.2)$$

for any $x \in \mathbb{R}$ and L^2-pricing measure Q. For any such Q with $\mathrm{Var}[Z_T] \leq \delta_{SR}^2$, there are hence no classical (Q-)good deals of any level $\delta > \delta_{SR}$. In other words, "too good" deals in terms of Sharpe ratios are excluded by an upper bound on the pricing kernel variance.

Remark 2.2 When choosing a number for the bound on the pricing kernel variance, one can think of two approaches. From empirical literature, one may have a quantitative idea of what "too large a Sharpe ratio" means. An alternative from theory is to observe that by (2.2), the smallest number one can choose is $\inf \sqrt{\mathrm{Var}[Z_T]}$ with the infimum taken over all L^2-pricing measures Q. This is typically attained by the so-called *variance-optimal martingale measure*, and the point is that thinking about specifying Sharpe ratio restrictions naturally brings up a dual optimisation problem over pricing measures. This will come up again at the end of this section and also motivates our subsequent Definition 3.2 of pricing kernel bounds.

Let us now measure the performance of X not by its Sharpe ratio, but by an expected utility $E[U(X)]$. In analogy to the classical approach, we want to define a *U-good deal* under Q as a payoff X whose expected utility $E[U(X)]$ is unduly large compared to the utility $U(E_Q[X])$ of its Q-price; and we want to exclude U-good deals via a bound on a suitable (U-dependent) functional of the pricing kernel. To that end, we need a duality result similar to (2.2). More precisely, given an initial capital x, we try to express the maximal expected utility from Q-affordable payoffs, $\sup_{X \in \mathcal{C}(x,Q)} E[U(X)]$, via the utility, $U(x)$, of the initial capital and some quantity, $f(Q|P)$, depending only on the pricing kernel. (The role of x in (2.2) becomes clearer if one replaces the Sharpe ratio with a quadratic "utility"; see (2.9) below). Imposing an upper bound A on $f(Q|P)$ then yields a set $\mathcal{N}(A)$ of no-good-deal measures Q and again rules out "too good" deals, measured now in terms of $E[U(X)]$.

To recall the well-known duality results we need, we now specify our financial market. $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ is a filtration on $(\Omega, \mathcal{F}, P)$ satisfying the usual conditions and, for simplicity, with $\mathcal{F}_T = \mathcal{F}$ and $\mathcal{F}_0$ trivial. Our numeraire is the bank account, and the $\mathbb{R}^d$-valued RCLL $(\mathbb{F}, P)$-semimartingale $S = (S_t)_{0 \leq t \leq T}$ describes the (discounted)

price evolution of d primary traded assets. S is assumed locally bounded for ease of exposition; see Remark 2.5 for comments on this. Absence of arbitrage is ensured by imposing on S the condition NFLVR; equivalently, the set $\mathcal{M}_e(S)$ of equivalent local martingale measures for S is not empty. An *admissible trading strategy*, $H \in \mathcal{H}$ for short, is an $\mathbb{R}^d$-valued predictable S-integrable process $H = (H_t)_{0 \le t \le T}$ such that the stochastic integral process $H{\cdot}S = \int H\, dS$ is uniformly bounded from below. For initial capital $x \in \mathbb{R}$, the set of payoffs *affordable in the market* is then

$$\mathcal{C}(x, S) := \left\{ X \in L^0 \,\middle|\, X^- \in L^\infty \text{ and } X \le x + (H{\cdot}S)_T \text{ for some } H \in \mathcal{H} \right\},$$

where we again (naturally) only allow payoffs bounded from below. It is well known that

$$\mathcal{C}(x, S) = \bigcap_{Q \in \mathcal{M}_e(S)} \mathcal{C}(x, Q), \tag{2.3}$$

and therefore we use as *allowed pricing measures* all $Q \in \mathcal{M}_e(S)$.

A *utility function* $U : \mathbb{R} \to \mathbb{R} \cup \{-\infty\}$ is increasing, continuous on its domain dom $U := \{x \in \mathbb{R} \mid U(x) > -\infty\}$, C^1 and strictly concave on the interior of dom U, and satisfies, with $\ell = \inf(\text{dom } U)$, the Inada conditions $U'(\infty) = 0$ and $U'(\ell) = +\infty$. Our examples have $\ell = 0$ or $\ell = -\infty$. The corresponding indirect utility is

$$u(x) := \sup_{H \in \mathcal{H}} E\left[U\big(x + (H{\cdot}S)_T\big) \right] = \sup_{X \in \mathcal{C}(x, S)} E[U(X)] \text{ for } x \in \text{dom } U,$$

setting $\infty - \infty = -\infty$ here, and to avoid degenerate situations, we assume that

$$u(x) < U(\infty) \text{ for all } x \in \text{dom } U. \tag{2.4}$$

The conjugate function of U is $V(y) := \sup_{x \in \mathbb{R}} \big(U(x) - xy\big)$ for $y > 0$, and we also need

$$v(y) := \inf_{Q \in \mathcal{M}_e(S)} E[V(yZ_T)] \text{ for } y > 0;$$

recall that $Z_T = \frac{dQ}{dP}$. Under the above assumptions, we have, from Kramkov and Schachermayer (1999) for $\ell = 0$ and Schachermayer (2001) for $\ell = -\infty$, the well-known duality relation

$$u(x) = \inf_{y > 0} \big(v(y) + xy\big) \text{ for } x \in \text{dom } U. \tag{2.5}$$

We need no asymptotic elasticity condition since we do not assert that v is finite everywhere.

For a utility analogue of (2.2), we now fix a pricing measure Q in $\mathcal{M}_e(S)$ and define

$$
\begin{aligned}
u^Q(x) &:= \sup_{X \in \mathcal{C}(x, Q)} E[U(X)] &&\text{for } x \in \text{dom } U, \\
v^Q(y) &:= E[V(yZ_T)] &&\text{for } y > 0.
\end{aligned}
$$

Arguments as for (2.5), or Theorem 10 of Frittelli (2000a), then yield the following *key duality*.

Proposition 2.3 *Suppose that $Q \in \mathcal{M}_e(S)$ satisfies the analogue of* (2.4), *namely*

$$u^Q(x) < U(\infty) \text{ for all } x \in \text{dom } U. \tag{2.6}$$

Then

$$u^Q(x) = \inf_{y>0} \left(v^Q(y) + xy \right) \text{ for } x \in \text{dom } U. \tag{2.7}$$

The next result shows that Proposition 2.3 is not vacuous.

Lemma 2.4 *Under the assumption* (2.4), *the set* $\mathcal{Q}_U := \{Q \in \mathcal{M}_e(S) \mid Q \text{ satisfies } (2.6)\}$ *is not empty. Moreover, we then have*

$$u(x) = \inf_{y>0} \left(\inf_{Q \in \mathcal{M}_e(S)} v^Q(y) + xy \right) = \inf_{Q \in \mathcal{M}_e(S)} u^Q(x). \tag{2.8}$$

Proof: By using (2.5), the definitions of v and v^Q, and (2.7), we directly obtain (2.8). If $Q \in \mathcal{M}_e(S)$ is not in $\mathcal{Q}_U$, then $u^Q(x) \geq U(\infty) > u(x)$ by (2.4), and so (2.8) implies that $u(x) = \inf_{Q \in \mathcal{Q}_U} u^Q(x)$. But if $\mathcal{Q}_U = \emptyset$, this would yield $u(x) = \infty$, contradicting (2.4). $\square$

Remark 2.5 We assume S locally bounded to use the results in Schachermayer (2001), which we need to include utility functions U with dom $U = \mathbb{R}$. The recent work in Biagini and Frittelli (2006) shows that one can also deal with unbounded processes; but the required technicalities would make the presentation unduly complicated for our purposes here.

The duality (2.7) in Proposition 2.3 is the general utility analogue of (2.2) which linked the maximal Sharpe ratio to the pricing kernel variance. But (2.7) is still too abstract, since it does not show clearly how the maximal expected utility $u^Q(x)$ is related to the utility $U(x)$ of the initial capital. To see what we need, consider for fixed $a \in \mathbb{R}$ the quadratic "utility" function $x \mapsto U^q(x) := -(a-x)^2$ on $\mathbb{R}$. One can show directly that (2.7) here takes the form

$$u^{q,Q}(x) := \sup_{X \in \mathcal{C}(x,Q)} E[U^q(X)] = U^q(x) \big/ E[Z_T^2] \text{ for } x < a. \tag{2.9}$$

Even though U^q is not a good utility function, (2.9) is useful. For one thing, it shows that U^q-good deals are excluded by a bound on the variance of the pricing kernel Z_T, and this approach is even equivalent to the classical one via Sharpe ratios. But more importantly, (2.9) explicitly splits $u^{q,Q}(x)$ into the "utility" $U^q(x)$ of the initial capital x and a term (here an L^2-norm) involving only the pricing measure Q. In order to obtain this for our U as well, we focus on a number of specific utility functions which allow more explicit computations.

The three examples we study in the sequel are

exponential utility: $U^e(x) = -e^{-\alpha x}/\alpha$ for $x \in \mathbb{R}$, with some $\alpha > 0$;
power utility: $U^p(x) = x^{1-\gamma}/(1-\gamma)$ for $x > 0$, with some $\gamma > 0, \gamma \neq 1$;
logarithmic utility: $U^\ell(x) = \log x$ for $x > 0$.

For each $i \in \{e, p, \ell\}$, we compute first the conjugate function $V^i(y)$, then the dual value function $v^{i,Q}(y) = E[V^i(yZ_T)]$, and finally use the key duality (2.7). Either in this way or by looking up Section 4.1 of Frittelli (2000a) or Section 5 of Bellini and Frittelli (2002), we obtain

Proposition 2.6 *For exponential, power and logarithmic utilities, (2.7) specialises to*

$$u^{e,Q}(x) = \sup_{X \in \mathcal{C}(x,Q)} E[U^e(X)] = U^e(x)e^{-f^e(Q|P)}, \tag{2.10}$$

$$u^{p,Q}(x) = \sup_{X \in \mathcal{C}(x,Q)} E[U^p(X)] = U^p(x)\left(\tfrac{1-\gamma}{\gamma}f^p(Q|P)\right)^\gamma, \tag{2.11}$$

$$u^{\ell,Q}(x) = \sup_{X \in \mathcal{C}(x,Q)} E[U^\ell(X)] = U^\ell(x) + f^\ell(Q|P), \tag{2.12}$$

respectively. In (2.10)–(2.12), we have used the divergences

$$f^e(Q|P) := E[Z_T \log Z_T] =: I(Q|P) \quad \text{(relative entropy)},$$
$$f^p(Q|P) := \tfrac{\gamma}{1-\gamma}E\left[Z_T^{\frac{1-\gamma}{\gamma}}\right],$$
$$f^\ell(Q|P) := E[-\log Z_T] = I(P|Q) \quad \text{(reverse relative entropy)}.$$

We are now ready for defining good deals.

Definition 2.7 *Fix $\delta > 0$, $i \in \{e, p, \ell\}$ and a pricing measure*

$$Q \in \mathcal{Q}_{U^i} := \left\{Q \in \mathcal{M}_e(S) \,\big|\, Q \text{ satisfies (2.6) with } u^{i,Q} \text{ and } U^i\right\}.$$

A payoff $X \in \bigcup_{x \in \mathrm{dom}\, U^i} \mathcal{C}(x, Q)$ is called a U^i-good deal of level δ with respect to Q if

a) $i = e$ and $E[U^e(X)] \geq \delta^{-1}U^e(E_Q[X])$, or

b) $i = p$ and $E[U^p(X)] \geq \delta^{\mathrm{sign}(1-\gamma)}U^p(E_Q[X])$, or

c) $i = \ell$ and $E[U^\ell(X)] \geq U^\ell(E_Q[X]) + \delta$.

For a constant $A \in \mathbb{R}$, the set of no-U^i-good-deal measures corresponding to A is

$$\mathcal{N}^i(A) := \left\{Q \in \mathcal{M}_e(S) \,\big|\, f^i(Q|P) \leq A\right\}. \tag{2.13}$$

Remark 2.8 **1)** The definition of U^i-good deals ensures by design that larger δ lead to good deals with respect to a higher utility level, because U^e and U^p for $\gamma > 1$ are nonpositive.

2) The divergences $f^i(Q|P)$ quantify the distance from Q to P since they take their minimal value if and only if $Q = P$; see Liese and Vajda (1987) and Goll and Rüschendorf (2001). Thus $Q \in \mathcal{N}^i(A)$ also means that the pricing measure Q is not too far away from the original measure P, in a sense quantified by the level A. This interpretation needs some care: For $i = p$ with $\gamma > 1$, we have $f^p(Q|P) \leq 0$ so that only values $A < 0$ make sense there.

While the above duality results are well known, we see our new contribution at a different level. We provide a general definition of U-good deals, and show how a systematic use of duality leads to clean results on a general probability space. In particular, our approach produces automatically from the chosen utility function U the appropriate functional $f(Q|P)$ on pricing kernels which in turn induces the set of no-good-deal measures. In our view, this contrasts favourably with the approach in Černý (2003) where rigorous results are only given in Section 2 for finite Ω. The derivation of generalised Sharpe ratios on pricing kernels in Section 3 of Černý (2003) uses Taylor expansions, asymptotics, transformations and scaling arguments, and we find many of these hard to follow.

Definition 2.7 and Proposition 2.6 immediately imply that a bound on the pricing kernel $\frac{dQ}{dP}$ excludes "too good" deals with respect to Q. Indeed, straightforward computations give

Corollary 2.9 *Take $i \in \{e, p, \ell\}$ and define the bounds $\delta^i_*(A)$ by*

a) $\delta^e_*(A) = e^A$ for $i = e$;

b) $\delta^p_*(A) = \left(\frac{1-\gamma}{\gamma} A\right)^{\gamma \, \text{sign}(1-\gamma)}$ for $i = p$;
   note that this also requires that $\text{sign} \, A = \text{sign}(1 - \gamma)$;

c) $\delta^\ell_*(A) = A$ for $i = \ell$.

Then there are no U^i-good deals of any level $\delta > \delta^i_*(A)$ with respect to any $Q \in \mathcal{N}^i(A) \cap \mathcal{Q}_{U^i}$.

A closer examination shows that the intersection with $\mathcal{Q}_{U^i}$ in Corollary 2.9 is not necessary; see the proof of Proposition 3.2.10 in Klöppel (2006). So like in the introduction, the no-good-deal value range for any payoff X is the interval from $\inf\limits_{Q \in \mathcal{N}^i(A)} E_Q[X]$ to $\sup\limits_{Q \in \mathcal{N}^i(A)} E_Q[X]$. To complete our approach by determining a set $\mathcal{N}^i(A)$ of no-U^i-good-deal measures, it remains to choose A in (2.13), and this has both a conceptual and a practical side. To ensure $\mathcal{N}^i(A) \neq \emptyset$, we must by (2.13) take

$$A > A^i_* := \inf_{Q \in \mathcal{M}_e(S)} f^i(Q|P),$$

which automatically leads us to consider the f^i-*optimal measure*

$$Q^i_* := \arg\min\left\{ f^i(Q|P) \, \middle| \, Q \in \mathcal{M}_e(S) \right\},$$

if it exists. (Of course, this always holds if Ω is finite.) This Q^i_* serves as a reference measure for U^i-good deals, in that its f^i-divergence $f^i(Q^i_*|P) = A^i_*$ is the smallest upper bound we can impose on pricing kernels to still obtain a nonempty set of no-good-deal measures. Moreover, A^i_* also has a nice economic interpretation. In fact, a look at

(2.10)–(2.12) shows that computing A^i_* amounts, up to U^i-dependent transformations, to computing the infimum over $Q \in \mathcal{M}_e(S)$ of $u^{i,Q}(x)$. But (2.8) in Lemma 2.4 gives

$$\inf_{Q \in \mathcal{M}_e(S)} u^Q(x) = u(x) = \sup_{X \in \mathcal{C}(x,S)} E[U(X)], \qquad (2.14)$$

and so *the smallest bound yielding a nonempty set of no-good-deal measures is directly linked to the maximal expected utility achievable by investing in the market.*

On the practical side, how should one choose a number for $A > A^i_*$? Since our measurement of good deals is too recent to have been analyzed empirically, we offer the following suggestion. Sharpe ratios have been extensively studied, and there is some agreement on when they are "too big", i.e., exceed some critical value b_{SR}. By (2.2), this translates into a bound of b^2_{SR} on the pricing kernel variance $\mathrm{Var}[Z_T]$. If we now work with a divergence $f(Q|P) = E\left[\phi\left(\frac{dQ}{dP}\right)\right]$, we can replace ϕ by its second order Taylor approximation to obtain an approximation of $f(Q|P)$ in terms of $\mathrm{Var}[Z_T]$, and we can translate any bound b^2_{SR} into an approximate bound for $f(Q|P)$. For our utility functions U^i, this yields

$$f^e(Q|P) = I(Q|P) \approx \tfrac{1}{2}\,\mathrm{Var}[Z_T],$$

$$f^p(Q|P) = \tfrac{\gamma}{1-\gamma} E\left[Z_T^{\frac{1-\gamma}{\gamma}}\right] \approx 1 + \tfrac{1-2\gamma}{2\gamma}\,\mathrm{Var}[Z_T]$$

$$f^\ell(Q|P) = I(P|Q) \approx \tfrac{1}{2}\,\mathrm{Var}[Z_T].$$

This approximation looks similar to Černý (2003), but is in a very different spirit: We use it to bridge the gap between theory and practical choices, but our results do not rely on it. We also remark that at least approximately, we obtain from the above that

$$f^e(Q|P) \approx f^\ell(Q|P) \approx \frac{\gamma}{1-2\gamma}\big(f^p(Q|P) - 1\big).$$

This could be used to relate or compare the bounds employed for the different utility functions.

3 Dynamic good deal bounds

In the static case, the *main steps of our approach to no-good-deal valuation* are:

1) Choose a (von Neumann–Morgenstern) utility function U to describe preferences.

2) Determine via (2.7) the corresponding divergence functional f on pricing measures Q.

3) For a suitable bound A, choose the set of no-good-deal pricing measures as

$$\mathcal{N}(A) := \big\{ Q \in \mathcal{M}_e(S) \,\big|\, f(Q|P) \leq A \big\}.$$

This precludes the existence of U-good deals of any level $\delta > \delta_*(A)$ under any $Q \in \mathcal{N}(A)$.

4) The value range for any payoff X is the interval from $\displaystyle\inf_{Q \in \mathcal{N}(A)} E_Q[X]$ to $\displaystyle\sup_{Q \in \mathcal{N}(A)} E_Q[X]$.

To extend this to a dynamic setting, we use the same financial market $(\Omega, \mathcal{F}, \mathbb{F}, P, S)$ as in Section 2 and define for time-T payoffs X upper and lower valuation bounds at any time t as essential supremum and infimum of $E_Q[X|\mathcal{F}_t]$, with Q ranging through a set of no-good-deal measures. The new issue in the dynamic context is to choose that set, and the message of this section is that naive generalisations from the static to the dynamic case do not yield nice valuation bounds over time. The key idea is therefore to use the approach from Section 2 not globally, but *locally, at the level of instantaneous market prices of risk.* As we shall see more clearly below, this is better than both the analogy generalisation $\mathcal{N}_t(A_t)$ from the static to the conditional setting and the lazy choice $\mathcal{N}_t(A) \equiv \mathcal{N}_0(A)$.

For ease of presentation, we focus in this section on exponential utility $U(x) = -e^{-\alpha x}/\alpha$. Comments on other utility functions appear in Section 5. S need no longer be locally bounded, but we still exclude arbitrage by assuming $\mathcal{M}_e(S) \neq \emptyset$; see later for comments. To derive for any time t an $\mathcal{F}_t$-conditional version of the duality (2.10), we introduce for initial capital $x_t \in L^\infty(\mathcal{F}_t)$ the set

$$\mathcal{C}_t(x_t, Q) := \left\{ X \in L^0 \,\middle|\, X^- \in L^\infty \text{ and } E_Q[X|\mathcal{F}_t] \leq x_t \right\}$$

of Q-affordable payoffs and set $u_t^Q(x_t) := \operatorname*{ess\,sup}_{X \in \mathcal{C}_t(x_t, Q)} E_Q[U(X)|\mathcal{F}_t]$. With the density process $(Z_t)_{0 \leq t \leq T}$ of Q with respect to P, the conditional analogue of $f(Q|P) = I(Q|P) = E[Z_T \log Z_T]$ is $f_t(Q|P) := E\left[\frac{Z_T}{Z_t} \log \frac{Z_T}{Z_t} \,\middle|\, \mathcal{F}_t \right]$, and one easily checks that each $f_t(Q|P)$ is finite if $f(Q|P)$ is. We assume

$$\mathcal{Q}^e := \left\{ Q \in \mathcal{M}_e(S) \,\middle|\, f(Q|P) < \infty \right\} \neq \emptyset,$$

and we call the f^e-optimal measure $Q_*^e := Q^E := \arg\min \left\{ f(Q|P) \,\middle|\, Q \in \mathcal{Q}^e \right\}$ the *minimal entropy martingale measure (MEMM)* for S. Q^E is always unique, and it exists by Theorem 2.1 of Frittelli (2000b) if S is locally bounded. For general S, existence and/or construction of Q^E are model-dependent. (This is a minor technical difference between general and finite Ω; in the latter case, infima are of course always minima.) The conditional analogue of the duality in Proposition 2.6 is

Proposition 3.1 *Assume $\mathcal{Q}^e \neq \emptyset$. For any $x_t \in L^\infty(\mathcal{F}_t)$ and $Q \in \mathcal{Q}^e$, we then have*

$$u_t^Q(x_t) = \operatorname*{ess\,sup}_{X \in \mathcal{C}_t(x_t, Q)} E[U(X)|\mathcal{F}_t] = U(x_t)e^{-f_t(Q|P)}. \tag{3.1}$$

Proof: Since $U(x) \leq V(y) + xy$ for $x \in \mathbb{R}$, $y > 0$, we get by choosing $y := \frac{Z_T}{Z_t} \exp\left(-\alpha x_t - f_t(Q|P) \right)$, conditioning on $\mathcal{F}_t$ and using $X \in \mathcal{C}_t(x_t, Q)$ and the explicit expression for V gives

$$E[U(X)|\mathcal{F}_t] \leq U(x_t)e^{-f_t(Q|P)}.$$

This proves "$\leq$" in (3.1). To obtain equality, take $X := -\frac{1}{\alpha}\left(\log \frac{Z_T}{Z_t} - \alpha x_t - f_t(Q|P) \right)$. If $X^- \in L^\infty$, then $X \in \mathcal{C}_t(x_t, Q)$ and "$\geq$" follows. Otherwise set $A_n := \{Z_T, Z_t \leq n\}$ and

$$X_n := -\frac{1}{\alpha}\left(I_{A_n} \log \frac{Z_T}{Z_t} - \alpha x_t - E\left[I_{A_n} \frac{Z_T}{Z_t} \log \frac{Z_T}{Z_t} \,\middle|\, \mathcal{F}_t \right] \right)$$

so that $X_n^- \in L^\infty$ and $E_Q[X_n|\mathcal{F}_t] = x_t$. Because $x \mapsto x \log x$ is bounded from below, monotone convergence first gives

$$\lim_{n \to \infty} E\left[I_{A_n} \tfrac{Z_T}{Z_t} \log \tfrac{Z_T}{Z_t} \,\Big|\, \mathcal{F}_t\right] = f_t(Q|P).$$

Since U is bounded from above, we then get $\lim_{n \to \infty} E[U(X_n)|\mathcal{F}_t] = E[U(X)|\mathcal{F}_t]$, again by monotone convergence and taking out $\mathcal{F}_t$-measurable factors, and "$\geq$" in (3.1) follows.
□

A conditional analogue to Section 2 would now be to choose $\mathcal{N}_t$ by imposing a bound A_t on $f_t(Q|P)$. However, we shall obtain much better dynamic properties for no-good-deal valuation bounds from conditions on the temporal derivative of $f_t(Q|P)$ or, equivalently, on the instantaneous market prices of risk. Making this precise requires more structure, and so we focus henceforth on a *Lévy model*.

So let $L = (L_t)_{0 \leq t \leq T}$ be an $\mathbb{R}^n$-valued Lévy process under P and $\mathbb{F}^L$ the P-augmentation of the filtration generated by L. All we need on Lévy processes, including notations used below, is summarised in the Appendix. Our results use the condition

$$\mathbb{F} = \mathbb{F}^L, \text{ i.e., the filtration is generated by a Lévy process.} \qquad (3.2)$$

This generalises the assumption, made explicitly or implicitly in almost all earlier papers, that $\mathbb{F}$ is generated by a Brownian motion (plus perhaps an independent Poisson process). It gives us a nice parametrisation for all $Q \approx P$ and a nice expression for $f_t(Q|P)$. Some aspects simplify if we also assume that $S = S_0 \, \mathcal{E}(L)$ is an exponential Lévy process, but we emphasise that our results do not need this. In particular, our setup includes certain stochastic volatility models with Lévy noise, as in Example 2 of Esche and Schweizer (2005).

Assume now $\mathbb{F} = \mathbb{F}^L$. Then any $Q \approx P$ has a density process of the form

$$Z = \mathcal{E}\left(\int \beta^{\mathrm{tr}} \, dL^c + \int \int_{\mathbb{R}^n} \big(Y(s,x) - 1\big)\big(\mu^L(dx, ds) - K(dx) \, ds\big)\right);$$

see Proposition A.3 in the Appendix. The processes β, Y are called *Girsanov parameters*, and we sometimes write $Q = Q^{(\beta, Y)}$ for emphasis. L is a Lévy process under $Q^{(\beta, Y)}$ if and only if β, Y are deterministic and time-independent; we then write $Q^{(\beta, Y)} = Q \in \mathcal{Q}^{\mathrm{Lévy}}$ and call β (then a constant) and $Y(\cdot)$ (then a function of $x \in \mathbb{R}^n$) time-independent for short. (Since $\mathcal{F}_0 = \mathcal{F}_0^L$ is trivial, time-independent does imply deterministic.) If $Q \in \mathcal{M}_e(S)$, i.e., Q is an ELMM for S, the Girsanov parameters have a very natural interpretation: β is the (instantaneous) *market price of diffusion risk*, and $Y - 1$ is the (instantaneous) *market price of jump risk*; see Appendix A of Björk and Slinko (2006). For any $Q \approx P$ with $f(Q|P) < \infty$, we can express $f_t(Q|P)$ explicitly in terms of β, Y and the Lévy triplet (b, c, K) of L under P. In fact, setting

$$g(y) := y \log y - y + 1 \geq 0 \text{ for } y \geq 0,$$
$$k\big(\beta, Y(\cdot)\big) := \tfrac{1}{2}\beta^{\mathrm{tr}} c\beta + \int_{\mathbb{R}^n} g\big(Y(x)\big) K(dx) \geq 0$$

for $\beta \in \mathbb{R}^n$ and a measurable function $Y : \mathbb{R}^n \to [0, \infty)$, we have

$$f_t\big(Q^{(\beta,Y)}\big|P\big) = E_{Q^{(\beta,Y)}}\left[\int_t^T k\big(\beta_s, Y_s(\cdot)\big)\, ds \,\middle|\, \mathcal{F}_t\right] \tag{3.3}$$

by an easy extension of Lemma 12 of Esche and Schweizer (2005). Note in (3.3) that β, Y appear twice—both in the integrand and in the conditioning measure. For $Q \in \mathcal{Q}^{\text{Lévy}}$ where β, Y are time-independent, (3.3) reduces to the non-random quantity $(T - t)\, k\big(\beta, Y(\cdot)\big)$.

Like in Section 2, we now define no-good-deal measures by imposing a bound on $f(Q|P)$ in a suitable way. We introduce a reference or "benchmark" measure $\widehat{Q} = Q^{(\hat{\beta},\hat{Y})}$ which is "close to the f-optimal ELMM Q^E" in the sense that $f(\widehat{Q}|P) \leq \inf_{Q \in \mathcal{Q}^e} f(Q|P) + \epsilon$. Note that such $\widehat{Q}$ exist as soon as $\mathcal{Q}^e \neq \emptyset$, for arbitrary $\mathbb{F}$ and whether or not the MEMM Q^E exists. We also point out that by Proposition 4.1 of Kabanov and Stricker (2002), Q^E (if it exists) is also $\mathcal{F}_t$-conditionally f-optimal for all t, i.e., $f_t\big(Q^E\big|P\big) \leq f_t(Q|P)$ for all $Q \in \mathcal{Q}^e$. (This also holds for power or logarithmic utilities, f^p or f^ℓ.)

Definition 3.2 *Assume $\mathbb{F} = \mathbb{F}^L$ and choose a "benchmark" measure $\widehat{Q} \in \mathcal{Q}^e$ and "close to Q^E". For predictable processes $\eta, \eta' \geq 1$ and $\theta, \theta' \geq 0$, we define*

$$\mathcal{N}_t^{\text{loc}}(\widehat{Q}) := \Big\{Q = Q^{(\beta,Y)} \in \mathcal{M}_e(S) \,\Big|\, k\big(\beta_s, Y_s(\cdot)\big) \leq \eta_s k\big(\hat{\beta}_s, \hat{Y}_s(\cdot)\big) + \theta_s$$

$$dP \otimes ds\text{-a.e. on } \Omega \times (t, T]\Big\}, \tag{3.4}$$

$$\mathcal{N}_t^{\text{glob}}(\widehat{Q}) := \Big\{Q = Q^{(\beta,Y)} \in \mathcal{M}_e(S) \,\Big|\, f_t(Q|P) \leq \eta_t' f_t(\widehat{Q}|P) + \theta_t' \ P\text{-a.s.}\Big\}.$$

We usually suppress the dependence of $\mathcal{N}_t^{\text{loc/glob}}(\widehat{Q})$ on $\eta, \theta, \eta', \theta'$ and call the restrictions on $f_t(Q|P)$ or Q the *local* or *global restrictions*, respectively, and $\mathcal{N}_t^{\text{loc/glob}}(\widehat{Q})$ the set of *local, global no-good-deal measures* at time t. Compared to Section 2, both restrictions should also be called *conditional*.

Remark 3.3   1) Our setting generalises both Černý (2003) by allowing jumps and Björk and Slinko (2006) by considering other than quadratic restrictions. However, our restrictions are again directly formulated in terms of instantaneous market prices of risk β, Y.
2) One may wonder why we include the parameter process θ or θ'. If $\theta' \equiv 0$, the upper bound for $f_t(Q|P)$ is proportional to $f_t(\widehat{Q}|P)$, and if $\widehat{Q}$ is chosen as the MEMM Q^E, then $f_t(\widehat{Q}|P)$ quantifies how far P is from being an ELMM for S. Now if P itself is already in $\mathcal{M}_e(S)$, this gives $\widehat{Q} = Q^E = P$ and so $\mathcal{N}_t^{\text{loc/glob}}(\widehat{Q})$ reduce to the singleton $\{P\}$. Of course one can argue that a subjective measure $P \in \mathcal{M}_e(S)$ is already a valid pricing measure; but it is still a matter of taste whether or not it is the only reasonable one. Allowing nonzero θ' or θ gives some extra freedom in that respect.

Proposition 3.4 *Suppose that* $\mathbb{F} = \mathbb{F}^L$ *and* $\mathcal{Q}^e \neq \emptyset$. *Then:*

1) *For any* $\widehat{Q} = Q^{(\hat{\beta}, \hat{Y})} \in \mathcal{Q}^e$, *we have* $\mathcal{N}_t^{\mathrm{loc}}(\widehat{Q}) \neq \emptyset$ *and* $\mathcal{N}_t^{\mathrm{glob}}(\widehat{Q}) \neq \emptyset$.

2) *The no-good-deal bounds obtained from* $\mathcal{N}_t^{\mathrm{loc}}$ *and from* $\mathcal{N}_0^{\mathrm{loc}}$ *coincide: For* $\widehat{Q} \in \mathcal{Q}^e$,

$$\operatorname*{ess\,inf}_{Q \in \mathcal{N}_t^{\mathrm{loc}}(\widehat{Q})} E_Q[X|\mathcal{F}_t] = \operatorname*{ess\,inf}_{Q \in \mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})} E_Q[X|\mathcal{F}_t] \text{ for any } X \in L^{\infty}. \tag{3.5}$$

The same is true for ess sup *instead of* ess inf.

Proof: 1) Clearly, $\mathcal{N}_t^{\mathrm{loc}}(\widehat{Q})$ and $\mathcal{N}_t^{\mathrm{glob}}(\widehat{Q})$ both contain $\widehat{Q}$.
2) By the Bayes rule, $\mathcal{F}_t$-conditional expectations under Q depend on Q only via the ratio Z_T/Z_t, which in turn depends only on the values of the Girsanov parameters β, Y on $]\!]t, T]\!]$; see Appendix, (A.3) in Proposition A.3. Hence restricting β and Y as in (3.4) has the same effect whether it is done for $(t, T]$ or for $(0, T]$. Indeed, "$\leq$" in (3.5) is obvious, and "$\geq$" follows by taking any $Q = Q^{(\beta, Y)} \in \mathcal{N}_t^{\mathrm{loc}}(\widehat{Q})$ and defining $Q' \in \mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})$ with $(\beta', Y') = (\beta, Y)$ on $\Omega \times (t, T]$ by $(\beta', Y') := (\hat{\beta}, \hat{Y}) I_{\Omega \times [0, t]} + (\beta, Y) I_{\Omega \times (t, T]}$. $\qquad\square$

Before our next result, we need two more definitions. We denote by $\mathcal{P}^{\ll}$ the set of all probability measures $Q \ll P$.

Definition 3.5 *A set* $\mathcal{S} \subseteq \mathcal{P}^{\ll}$ *with* $\mathcal{S} \cap \mathcal{P}^{\approx} \neq \emptyset$ *is called* m-stable *if it has the following property: Whenever we take* $Q^1, Q^2 \in \mathcal{S}$ *with density processes* Z^1, Z^2 *with respect to* P, *fix a stopping time* $\tau \leq T$ *and impose* $Q^2 \approx P$, *then* $Z_T := Z_\tau^1 \frac{Z_T^2}{Z_\tau^2}$ *is the density of some* $Q \in \mathcal{S}$. *(This means that* $\mathcal{S}$ *is closed under "switching from* Q^1 *to* Q^2 *at a stopping time* τ".)

Definition 3.6 *A* monetary concave utility functional (MCUF) *at time* t *on* L^{∞} *is a mapping* $\Phi_t : L^{\infty} \to L^{\infty}(\mathcal{F}_t)$ *satisfying*

A) monotonicity: $X_1 \leq X_2$ *implies* $\Phi_t(X_1) \leq \Phi_t(X_2)$;

B) $\mathcal{F}_t$-translation invariance: $\Phi_t(X + a_t) = \Phi_t(X) + a_t$ *for* $a_t \in L^{\infty}(\mathcal{F}_t)$;

C) concavity: $\Phi_t\big(\alpha X_1 + (1 - \alpha)X_2\big) \geq \alpha\Phi_t(X_1) + (1 - \alpha)\Phi_t(X_2)$ *for* $\alpha \in [0, 1]$.

The acceptance set *of* Φ_t *is* $\mathcal{A}_t := \{X \in L^{\infty} \mid \Phi_t(X) \geq 0\}$. Φ_t *is a* monetary coherent utility functional (MCohUF) *at time* t *if it satisfies in addition*

D) positive homogeneity: $\Phi_t(\lambda X) = \lambda\Phi_t(X)$ *for* $\lambda \geq 0$.

In C) and D), we could equivalently postulate the properties for α and $\lambda \in L^{\infty}$ which are $\mathcal{F}_t$-measurable instead of constant. A *(static) convex/coherent risk measure* ρ on L^{∞} is a mapping such that $\Phi_0 = -\rho$ is an MCUF/MCohUF at time 0. Finally, a *dynamic MCUF (or MCohUF)* is a family $\Phi = (\Phi_t)_{0 \leq t \leq T}$ where each Φ_t is an MCUF (or MCohUF) at time t. Risk measures and/or monetary utility functionals have attracted a lot of attention recently. Chapter 4 of Föllmer and Schied (2004) gives a textbook account on static risk

measures, and Klöppel and Schweizer (2005) contains a partial overview of recent work on dynamic MCUFs. See also Section 4 for a discussion on the link to abstract good deal bounds.

Fix a benchmark measure $\widehat{Q} \in \mathcal{Q}^e$ (mostly suppressed in the notation) and predictable $\eta \geq 1$, $\theta \geq 0$, and define *lower and upper no-good-deal valuation bounds* for $X \in L^\infty$ by

$$
\begin{aligned}
\pi_t^\ell(X) &:= \operatorname*{ess\,inf}_{Q \in \mathcal{N}_t^{\mathrm{loc}}(\widehat{Q})} E_Q[X|\mathcal{F}_t], \\
\pi_t^u(X) &:= \operatorname*{ess\,sup}_{Q \in \mathcal{N}_t^{\mathrm{loc}}(\widehat{Q})} E_Q[X|\mathcal{F}_t] = -\pi_t^\ell(-X).
\end{aligned}
\tag{3.6}
$$

The first main result in this section shows that $\pi^\ell(\cdot)$ is a dynamic monetary coherent utility functional with a very good dynamic behaviour over time. Note that this result does not depend on the choice of the benchmark measure $\widehat{Q}$.

Theorem 3.7 *Assume $\mathbb{F} = \mathbb{F}^L$ and $\mathcal{Q}^e \neq \emptyset$, and fix $\widehat{Q} \in \mathcal{Q}^e$. Then:*

1) *The set $\mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})$ is m-stable.*

2) *$\pi^\ell = (\pi_t^\ell)_{0 \leq t \leq T}$ is a dynamic monetary coherent utility functional (DMCohUF). $\big($As an aside, this is also true if we use $\mathcal{N}_t^{\mathrm{glob}}(\widehat{Q})$ instead of $\mathcal{N}_t^{\mathrm{loc}}(\widehat{Q})$.$\big)$*

3) *For each $X \in L^\infty$, there exists an RCLL version of $\big(\pi_t^\ell(X)\big)_{0 \leq t \leq T}$ such that*

$$
\pi_\tau^\ell(X) = \operatorname*{ess\,inf}_{Q \in \mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})} E_Q[X|\mathcal{F}_\tau] \ P\text{-a.s. for each stopping time } \tau \leq T.
$$

4) *π^ℓ is (stopping-time-)consistent in the sense that for any stopping times $\sigma \leq \tau \leq T$, $\pi_\tau^\ell(X_1) \leq \pi_\tau^\ell(X_2)$ implies $\pi_\sigma^\ell(X_1) \leq \pi_\sigma^\ell(X_2)$.*

5) *π^ℓ is recursive (with respect to stopping times) in the sense that $\pi_\sigma^\ell\big(\pi_\tau^\ell(X)\big) = \pi_\sigma^\ell(X)$ for any stopping times $\sigma \leq \tau \leq T$.*

Proof: $\mathcal{M}_e(S)$ is m-stable by Proposition 5 of Delbaen (2006) (this does not need that S is locally bounded), and so is then clearly $\mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})$. Moreover, (3.6) easily yields that each π_t^ℓ is an MCohUF at time t. Indeed, this uses only that π_t^ℓ is an ess inf over conditional expectations and thus it still holds for $\mathcal{N}_t^{\mathrm{glob}}(\widehat{Q})$ instead of $\mathcal{N}_t^{\mathrm{loc}}(\widehat{Q})$. All other properties of π^ℓ stem from the m-stability of $\mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})$. First of all, by part 2) of Proposition 3.4, the ess inf in (3.6) can be taken over $\mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})$ instead of $\mathcal{N}_t^{\mathrm{loc}}(\widehat{Q})$. Since $\mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})$ is m-stable, Lemmas 22 and 23 in Delbaen (2006) now imply part 3), and Theorem 12 in Delbaen (2006) yields parts 4) and 5). We remark that the results we use here do not need the L^1-closedness assumption made in Delbaen (2006). $\square$

As the proof shows, the main work for Theorem 3.7 has been done in an abstract setting by Delbaen (2006). The importance of time-consistency and recursiveness is also discussed in Section 3 of Klöppel and Schweizer (2005) or Chapter 2 of Klöppel (2006).

One of our contributions here is to make clear why the choice of the set $\mathcal{N}$ of no-good-deal measures is so crucial. Indeed, as can be seen from the proof, the key to ensuring all the good temporal properties in Theorem 3.7 is the m-stability of $\mathcal{N}_0^{\mathrm{loc}}(\widehat{Q})$. This also allows to use dynamic programming techniques; see Section 8 of Delbaen (2006) for more details.

Theorem 3.7 shows that no-good-deal valuation bounds from conditional *local* restrictions on instantaneous market prices of risk are time-consistent. It is natural to ask if time-consistency also holds for bounds coming from *global* restrictions, and Example 3.8 shows that the answer is negative in general. More precisely, this example gives a setting where the (convex) L^1-closure of $\mathcal{N}_0^{\mathrm{glob}}(\widehat{Q})$ is not m-stable. By Lemma 2.3.29 of Klöppel (2006), this last condition is necessary for time-consistency of the no-good-deal valuation bounds obtained from $\mathcal{N}_0^{\mathrm{glob}}(\widehat{Q})$, and so if it fails, there must be payoffs for which time-consistency of the global bounds fails as well.

Example 3.8 Choose a model satisfying (3.2) and $S = S_0 \mathcal{E}(L)$ where $Q^E = Q^{(\beta_*, Y_*)}$ exists, and where there exists another $Q^{(\beta, Y)} \in \mathcal{Q}^e \cap \mathcal{Q}^{\mathrm{Lévy}}$ which is different from Q^E. By the discussion in a) – c) below, Q^E is then in $\mathcal{Q}^{\mathrm{Lévy}}$. Fix $t \in (0, T)$ and define Q^1, Q^2 via

$$(\beta^1, Y^1) := (\beta, Y)I_{[\![0,t]\!]} + (\beta_*, Y_*)I_{]\!]t,T]\!]},$$
$$(\beta^2, Y^2) := (\beta_*, Y_*)I_{[\![0,t]\!]} + (\beta, Y)I_{]\!]t,T]\!]}.$$

Then both Q^1, Q^2 are in $\mathcal{Q}^e$ and we readily compute from (3.3) that

$$\begin{aligned}
f(Q^1|P) &= t\,k(\beta, Y(\cdot)) + (T - t)k(\beta_*, Y_*(\cdot)), \\
f(Q^2|P) &= t\,k(\beta_*, Y_*(\cdot)) + (T - t)k(\beta, Y(\cdot)).
\end{aligned} \tag{3.7}$$

So if we set $\theta' \equiv 0$ and $\eta' \equiv \max\left(\frac{f(Q^1|P)}{f(Q^E|P)}, \frac{f(Q^2|P)}{f(Q^E|P)}\right)$, then also $Q^1, Q^2 \in \mathcal{N}_0^{\mathrm{glob}}(Q^E)$.

Now observe that by switching from Q^1 to Q^2 at time t, i.e., via $Z_T = Z_t^1 \frac{Z_T^2}{Z_t^2}$, we obtain $Q = Q^{(\beta, Y)}$. Because all measures have deterministic Girsanov parameters, we get

$$Tk(\beta_*, Y_*(\cdot)) = f\left(Q^E|P\right) < f(Q|P) = Tk(\beta, Y(\cdot))$$

since $Q \neq Q^E$, and therefore by (3.7)

$$\eta' f\left(Q^E|P\right) + \theta' = \max\left(f(Q^1|P), f(Q^2|P)\right) < Tk(\beta, Y(\cdot)) = f(Q|P).$$

Thus $Q \notin \mathcal{N}_0^{\mathrm{glob}}(Q^E)$, and Q is not even in the L^1-closure of $\mathcal{N}_0^{\mathrm{glob}}(Q^E)$. Hence both these sets are not m-stable.

Despite an absence of time-consistency, one might still feel that a (global, conditional) bound on $f_t(Q|P)$ looks more intuitive than a (local, conditional) bound as in (3.4). To give further support for our choice of local, not global restrictions, we now argue that

under one more assumption, "local restrictions imply global ones"; thus the nice economic interpretations of global restrictions are also valid for local ones. More precisely, this works if the benchmark measure $\widehat{Q}$ has deterministic Girsanov parameters, i.e., is in $\mathcal{Q}^{\text{Lévy}}$. So we need $\mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}} \neq \emptyset$, and we also want $\widehat{Q}$ "close to Q^E". For generality, we shall impose this as a condition, but we nevertheless explain briefly how it can be ensured. At this point, the condition that $S = S_0 \mathcal{E}(L)$ comes in handy, since it is sufficient for several useful properties. More precisely:

a) If $S = S_0 \mathcal{E}(L)$ (and as before $\mathbb{F} = \mathbb{F}^L$), then L is by Theorem A of Esche and Schweizer (2005) a Lévy process under Q^E, provided that Q^E exists and $\mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}} \neq \emptyset$; thus even $Q^E \in \mathcal{Q}^{\text{Lévy}}$ and so we could choose $\widehat{Q} = Q^E$. If $\mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}} \neq \emptyset$ but we do not know whether Q^E exists, we can still find for any $\epsilon > 0$ a Lévy benchmark measure $\widehat{Q} \in \mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}}$ "close to Q^E", i.e., with $f(\widehat{Q}|P) \leq \inf\limits_{Q \in \mathcal{Q}^e} f(Q|P) + \epsilon$; see Proposition 18 of Esche and Schweizer (2005).

b) If $S = S_0 \mathcal{E}(L)$ and $\mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}} \neq \emptyset$, the only open point in a) is whether Q^E exists. This can be checked in specific models via the results of Fujiwara and Miyahara (2003) or Theorem B of Esche and Schweizer (2005). Conversely, *assuming $S = S_0 \mathcal{E}(L)$ plus existence of Q^E* (obtained perhaps by some direct argument) *implies* $\mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}} \neq \emptyset$, due to a).

c) For unbounded S, the no-arbitrage condition NFLVR is in general weaker than the property $\mathcal{M}_e(S) \neq \emptyset$. But $S = S_0 \mathcal{E}(L)$ and $\mathbb{F} = \mathbb{F}^L$ imply by Theorem 2.7 of Kardaras (2007) that the two conditions are equivalent, so that we are then in a standard setting.

In summary, adding $S = S_0 \mathcal{E}(L)$ to our standard assumption (3.2) that $\mathbb{F} = \mathbb{F}^L$ gives a number of nice properties. But we emphasise again that $S = S_0 \mathcal{E}(L)$ is not really needed.

Now we are ready for the second main result in this section. We recall that $\mathcal{Q}^{\text{Lévy}}$ is the set of all $Q \approx P$ such that L is a Lévy process under Q and define for $\widehat{Q} = Q^{(\hat{\beta}, \hat{Y})}$

$$\mathcal{N}_t^{\text{loc/glob, Lévy}}(\widehat{Q}) := \mathcal{N}_t^{\text{loc/glob}}(\widehat{Q}) \cap \mathcal{Q}^{\text{Lévy}}.$$

If $\hat{\beta}, \hat{Y}$ are deterministic and time-independent, it makes sense to require this for η, θ as well.

Theorem 3.9 *Assume $\mathbb{F} = \mathbb{F}^L$ and $\mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}} \neq \emptyset$. Fix $\widehat{Q} \in \mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}}$ and choose $\eta(s) \geq 1$ and $\theta(s) \geq 0$ as deterministic measurable functions on $[0, T]$. Then:*

1) The local restrictions imply the global ones: If we set $\eta'(t) := \frac{1}{T-t} \int\limits_t^T \eta(s)\, ds$ and $\theta'(t) := \int\limits_t^T \theta(s)\, ds$, then for any $\widehat{Q} \in \mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}}$,

$$\mathcal{N}_t^{\text{loc}}(\widehat{Q}) \subseteq \mathcal{N}_t^{\text{glob}}(\widehat{Q}). \tag{3.8}$$

2) If we even assume $\eta(s) \equiv \eta \geq 1$ and $\theta(s) \equiv \theta \geq 0$, then for Lévy pricing measures, the local and global restrictions coincide: If $\eta'(t) = \eta$ and $\theta'(t) = (T - t)\theta$ as in part 1), then

$$\mathcal{N}_t^{\text{loc, Lévy}}(\widehat{Q}) = \mathcal{N}_t^{\text{glob, Lévy}}(\widehat{Q}) =: \mathcal{N}_t^{\text{Lévy}}(\widehat{Q}) = \mathcal{N}_0^{\text{Lévy}}(\widehat{Q}). \tag{3.9}$$

Moreover, all Lévy no-good-deal bounds coincide as well:

$$\underset{Q \in \mathcal{N}_t^{\text{Lévy}}(\widehat{Q})}{\text{ess inf}} E_Q[X|\mathcal{F}_t] = \underset{Q \in \mathcal{N}_0^{\text{Lévy}}(\widehat{Q})}{\text{ess inf}} E_Q[X|\mathcal{F}_t] \text{ for any } X \in L^\infty,$$

and the same is true for ess sup *instead of* ess inf.

Proof: 1) Recall that $\hat{\beta}$ and $\hat{Y}$ are deterministic and time-independent, since $\widehat{Q} \in \mathcal{Q}^{\text{Lévy}}$. Due to (3.3), $f_t(\widehat{Q}|P) = (T-t)k(\hat{\beta}, \hat{Y}(\cdot))$ is therefore deterministic, and so are η and θ. Hence (3.8) follows immediately from (3.3) by the definitions in (3.4).
2) For $Q \in \mathcal{Q}^e \cap \mathcal{Q}^{\text{Lévy}}$, (3.3) gives $f_t(Q|P) = (T-t)k(\beta, Y(\cdot))$, which directly implies the first equality in (3.9). The last follows easily: Since $\widehat{Q}$ is in $\mathcal{Q}^{\text{Lévy}}$ and η, θ do not depend on t, the restriction in (3.4) for $Q \in \mathcal{Q}^{\text{Lévy}}$ is the same for all t. The rest is obvious. $\square$

Remark 3.10 1) The implication in part 1) of Theorem 3.9 extends similar statements in Černý (2003) by giving a rigorous formulation and by covering processes with jumps as well.
2) Part 2) of Theorem 3.9 is simple and unsurprising, but still of interest. It shows that if one is very restrictive about the Lévy structure, one can obtain the same no-good-deal measures and valuation bounds equivalently via local or global restrictions. This is useful since computations (under some Q) are often much simpler if L has the Lévy property.
3) Example 3.8 with $\eta' \equiv \frac{f(Q|P)}{f(Q^E|P)}$ also shows that $\mathcal{N}_0^{\text{loc/glob, Lévy}}(Q^E)$ is not m-stable in general; indeed, Q^E and Q are both in $\mathcal{Q}^{\text{Lévy}}$ and Q^1 is obtained by switching from Q to Q^E at time t, but Q^1 is not in $\mathcal{Q}^{\text{Lévy}}$ since β^1, Y^1 depend on time. So the price to pay for simpler computations under Lévy bounds is again a lack of time-consistency.

We conclude this section with comments on related literature. With the exception of Björk and Slinko (2006), all papers we know on dynamic good deal bounds in continuous-time models, including Cochrane and Saá-Requejo (2000) and Černý (2003), use approximations or heuristics in some way. These rely crucially on a diffusion setting and therefore cannot deal with continuous-time processes having jumps. The only rigorous work so far on dynamic good deals is Björk and Slinko (2006), who study conditional local quadratic bounds on the instantaneous market price of risk when $\mathbb{F}$ is generated by Brownian motion and a marked point process. Their main goal is to extend Cochrane and Saá-Requejo (2000) by including jumps and giving a rigorous and simpler approach. By using a quadratic instead of an exponential utility function, we could obtain the same no-good-deal measures as in Björk and Slinko (2006); see Proposition 4.4.3 of Klöppel (2006).

Our approach here is more general than the literature in two ways. We allow processes with jumps and use a general (non-quadratic) utility, and it is the combination of these that makes our results significantly different. While Černý (2003) also allows general utilities, he mentions in Section 5.1 that these always lead to the same, quadratic local restrictions for Itô processes. His model assumptions are not clearly spelled out, but the

underlying reason why one obtains that conclusion is as follows. It is well known that in suitably nice Markovian diffusion settings, many f-optimal measures coincide with the minimal martingale measure $Q^{\min}$, and this holds in particular for the three utilities discussed in Section 2. Hence we can then choose $Q^{\min}$ for the benchmark measure $\widehat{Q}$. Next, the functional k in (3.3) reduces to a quadratic function in the absence of jumps, and so the local restrictions indeed boil down to quadratic conditions based on $\widehat{Q} = Q^{\min}$. However, this changes as soon as we have jumps, because then (3.3) includes a second term coming from the Lévy measure K, and it is also well known that the f-optimal measures (and hence our typical benchmark measures $\widehat{Q}$) genuinely depend, via f, on the chosen utility function.

All in all, it seems that a rigorous, systematic and comparatively general study like ours on local and global restrictions for no-good-deal valuation in settings with jumps has not been done so far.

4   Good deal bounds and risk measures

Up to now, we have explained how to construct static and dynamic no-good-deal valuation bounds from an expected utility criterion. There is another closely related approach via risk measures, and we briefly present these ideas here to do justice to the whole breadth of the literature. In particular, we point out some of the resulting strong parallels.

For brevity and since our main goal is to provide a comparison, we only explain good deal bounds from risk measures for the static case. A dynamic (or rather conditional) version of the definition can be found at the end of Section 2.6 in Klöppel (2006), and the presentation there also shows (on p. 82) how the resulting bounds can be represented as the convolution of two dynamic risk measures. This is useful since such convolutions have been studied for time-consistency; see for instance Cheridito and Kupper (2006) and Klöppel and Schweizer (2007). But we do not go into details here.

Let $\Psi = \Phi_0$ be a monetary concave utility functional (MCUF) at time 0 and recall from Definition 3.6 that $\rho = -\Psi$ is a convex risk measure on L^∞ and that the acceptance set of Ψ (or ρ) is $\mathcal{A} := \{X \in L^\infty \mid \Psi(X) \geq 0\}$. We can argue equivalently via Ψ or via $\mathcal{A}$ because by Proposition 4.6 of Föllmer and Schied (2004), Ψ and $\mathcal{A}$ uniquely determine each other via

$$\Psi(X) = \sup\{m \in \mathbb{R} \mid X - m \in \mathcal{A}\} = \sup\big((X - \mathcal{A}) \cap \mathbb{R}\big) \text{ for all } X \in L^\infty. \quad (4.1)$$

Although the first abstract papers on good deals worked with risk measures, we use here the formulation via monetary utility functionals since the connection with von Neumann–Morgenstern expected utility becomes then clearer. We take the same financial market as in Section 2 and set $\mathcal{C}_b(0, S) := \mathcal{C}(0, S) \cap L^\infty$. As in Jaschke and Küchler (2001), a payoff $X \in \mathcal{C}_b(0, S)$ is an $\mathcal{A}$-*good deal* if $X - \epsilon \in \mathcal{A}$ for some $\epsilon > 0$; in words, X is superreplicable at zero cost and remains acceptable for Ψ even if one subtracts from it a little money. The *lower no-$\mathcal{A}$-good-deal valuation bound* for a payoff $X \in L^\infty$ is then

$$\pi_\mathcal{A}(X) := \sup\{m \in \mathbb{R} \mid X - m + X' \in \mathcal{A} \text{ for some } X' \in \mathcal{C}_b(0, S)\}.$$

To see why this is a lower bound, suppose one could buy X for $y < \pi_\mathcal{A}(X)$. Then there are $\epsilon > 0$ and $X' \in \mathcal{C}_b(0, S)$ with $y + \epsilon \leq \pi_\mathcal{A}(X) - \epsilon$ and $X - (y + \epsilon) + X' \in \mathcal{A}$. So

buying X for y and superreplicating X' at zero cost yields the total payoff $X - y + X'$, which is an $\mathcal{A}$-good deal. Hence X is too cheap at y. Analogously, the upper no-good-deal bound is $-\pi_\mathcal{A}(-X)$.

Remark 4.1 Conceptually, no-good-deal valuation via risk measures (or MCUFs) rests on the same idea as the expected-utility-based method in Section 2; one wants to avoid having payoffs which are "too good" in a sense specified by some utility. In contrast to von Neumann–Morgenstern expected utilities, monetary utility functionals are translation-invariant, which explains why (like for the Sharpe ratio) the initial capital does not appear explicitly in the definition of $\mathcal{A}$-good deals. A second, minor difference is that "good deal" and "acceptable" are here less clearly quantitative; an analogue of the level δ from Section 2 is well hidden in the choice of the set $\mathcal{A}$.

The above approach to defining and obtaining good deal bounds via risk measures goes back to Jaschke and Küchler (2001) who in turn refer to an earlier (draft) version of Černý and Hodges (2002). To give due credit, many key ideas appear already in Hodges (1998). More precisely, Jaschke and Küchler (2001) work with *coherent* risk measures and also establish a good deal version of the fundamental theorem of asset pricing. Staum (2004) generalises this to *convex* risk measures; see also Section 6 of Klöppel and Schweizer (2005) for a temporally dynamic treatment. Similar results can be found in Cherny (2007a,b) who, like Jaschke and Küchler (2001), replaces our specific description $\mathcal{C}(0, S)$ of the set of attainable claims by an abstract cone to cover for example transaction costs as well. An earlier special version is Carr, Geman, and Madan (2001) where acceptable payoffs are defined via valuation and stress test measures.

Now let $\Psi : L^\infty \to \mathbb{R}$ be an arbitrary mapping. By Corollary 4.34 in Föllmer and Schied (2004), Ψ is an MCohUF and continuous from above $\big($in the sense that for any sequence $(X_n)_{n\in\mathbb{N}} \subseteq L^\infty$ decreasing to some $X \in L^\infty$, we have $\lim_{n\to\infty} \Psi(X_n) = \Psi(X)\big)$ if and only if Ψ can, for some set $\mathcal{Q} \subseteq \mathcal{P}^\ll = \{$probability measures $Q \ll P\}$, be represented as

$$\Psi(X) = \inf_{Q\in\mathcal{Q}} E_Q[X] \text{ for all } X \in L^\infty. \tag{4.2}$$

We call (any such) $\mathcal{Q}$ a set of *representing measures* for Ψ and note that $\mathcal{Q}$ can always be chosen convex. In particular, the mapping $X \mapsto \inf_{Q\in\mathcal{N}} E_Q[X]$ on L^∞ is an MCohUF for any nonempty set $\mathcal{N} \subseteq \mathcal{P}^\ll$.

One important insight in Section 2 was that the smallest bound yielding a nonempty set of no-U-good-deal measures is directly linked to the maximal expected utility achievable by investing in the market; see (2.14). The analogous result in this section is that *the lower no-$\mathcal{A}$-good-deal valuation bound for X is exactly the maximal monetary Ψ-utility achievable by trading in the market with random endowment X.* Moreover, the second equality in (4.3) below shows that we can also interpret the set $\mathcal{Q}$ of representing measures for Ψ almost as a set of no-$\mathcal{A}$-good-deal measures. Identifying any subset of $\mathcal{P}^\ll$ with the corresponding set in $L^1(P)$ of densities $\frac{dQ}{dP}$, the precise formulation is as follows.

Proposition 4.2 *Take an MCohUF Ψ represented via* (4.2) *by a convex set $\mathcal{Q} \subseteq \mathcal{P}^\ll$ and suppose that $\mathcal{Q}$ is relatively weakly compact in $L^1(P)$. $\big($By Corollary 4.35 of Föllmer and*

Schied (2004), this amounts to a continuity property of Ψ.) Suppose that $\mathcal{Q} \cap \mathcal{M}_e(S) \neq \emptyset$ and set $\mathcal{N} := \overline{\mathcal{Q}} \cap \mathcal{M}_e(S)$, where the bar $\overline{}$ denotes the closure in $L^1(P)$. Then

$$\pi_{\mathcal{A}}(X) = \sup_{X' \in \mathcal{C}_b(0,S)} \Psi(X + X') = \inf_{Q \in \mathcal{N}} E_Q[X] \text{ for all } X \in L^\infty. \tag{4.3}$$

Actually, the first equality in (4.3) does not need the compactness assumption on $\mathcal{Q}$.

Proof: 1) If we rewrite the definition of $\pi_{\mathcal{A}}(X)$ and use (4.1), we directly obtain

$$\pi_{\mathcal{A}}(X) = \sup_{X' \in \mathcal{C}_b(0,S)} \sup\{m \in \mathbb{R} \mid X - m + X' \in \mathcal{A}\} = \sup_{X' \in \mathcal{C}_b(0,S)} \Psi(X + X'),$$

since $\mathcal{A}$ is the acceptance set of Ψ. This gives the first equality in (4.3).

2) First of all, (4.2) gives $\Psi(X) = \inf_{Q \in \mathcal{Q}} E_Q[X] = \inf_{Q \in \overline{\mathcal{Q}}} E_Q[X]$ for any $X \in L^\infty$. Since $\overline{\mathcal{Q}}$ is a compact Hausdorff space and $\mathcal{C}_b(0,S)$ is convex, the minimax theorem implies that

$$\sup_{X' \in \mathcal{C}_b(0,S)} \Psi(X + X') = \sup_{X' \in \mathcal{C}_b(0,S)} \inf_{Q \in \overline{\mathcal{Q}}} E_Q[X + X']$$
$$= \inf_{Q \in \overline{\mathcal{Q}}} \sup_{X' \in \mathcal{C}_b(0,S)} E_Q[X + X']. \tag{4.4}$$

By the optional decomposition theorem, $Q \in \mathcal{P}^\approx$ is in $\mathcal{M}_e(S)$ if and only if $E_Q[X'] \leq 0$ for all $X' \in \mathcal{C}_b(0,S)$; compare (2.3). Using $\mathcal{M}_e(S) \neq \emptyset$ and an approximation argument, we obtain that $Q \in \mathcal{P}^\ll$ is in $\overline{\mathcal{M}_e(S)}$ if and only if $E_Q[X'] \leq 0$ for all $X' \in \mathcal{C}_b(0,S)$. Because $\mathcal{C}_b(0,S)$ is a cone and contains 0, we thus have for any $Q \in \mathcal{P}^\ll$ that

$$\sup_{X' \in \mathcal{C}_b(0,S)} E_Q[X + X'] = \begin{cases} E_Q[X] & \text{if } Q \in \overline{\mathcal{M}_e(S)}, \\ +\infty & \text{otherwise.} \end{cases}$$

Combining this first with (4.4) and then using $\mathcal{Q} \cap \mathcal{M}_e(S) \neq \emptyset$ plus the definition of $\mathcal{N}$ yields

$$\sup_{X' \in \mathcal{C}_b(0,S)} \Psi(X + X') = \inf_{Q \in \overline{\mathcal{Q}} \cap \overline{\mathcal{M}_e(S)}} E_Q[X] = \inf_{Q \in \overline{\mathcal{Q}} \cap \mathcal{M}_e(S)} E_Q[X] = \inf_{Q \in \mathcal{N}} E_Q[X].$$

Indeed, to see the second equality above, note that "$\leq$" is obvious, and "$\geq$" follows if we show that any $Q \in \overline{\mathcal{Q}} \cap \overline{\mathcal{M}_e(S)}$ is the L^1-limit of a sequence $(Q^n)_{n \in \mathbb{N}}$ in $\overline{\mathcal{Q}} \cap \mathcal{M}_e(S)$. To that end, take $Q^* \in \mathcal{Q} \cap \mathcal{M}_e(S)$ and set $Q^n := \frac{1}{n} Q^* + \left(1 - \frac{1}{n}\right) Q$. Then L^1-convergence to Q is clear, and so is $Q^n \in \overline{\mathcal{Q}} \cap \overline{\mathcal{M}_e(S)}$ since both Q and Q^* lie in that convex set. But it is easy to check that each Q^n is also a local martingale measure for S, and since Q^n is like Q^* equivalent to P, we get $Q^n \in \overline{\mathcal{Q}} \cap \mathcal{M}_e(S)$. This completes the proof. $\square$

Remark 4.3 Since the upper no-$\mathcal{A}$-good-deal valuation bound is $-\pi_{\mathcal{A}}(-X) = \sup_{Q \in \mathcal{N}} E_Q[X]$, Proposition 4.2 shows how a given MCohUF Ψ induces a set $\mathcal{N} \subseteq \mathcal{M}_e(S)$ of valuation measures which exactly yield all the no-$\mathcal{A}$-good-deal values corresponding to Ψ. Hence $\mathcal{N}$ can be viewed as the set of *no-$\mathcal{A}$-good-deal valuation measures induced by Ψ*, and it is explicitly obtained as $\mathcal{N} = \overline{\mathcal{Q}} \cap \mathcal{M}_e(S)$, where $\mathcal{Q}$ represents Ψ.

Conversely, suppose we already have some set of no-good-deal valuation measures; this could like in Section 2 be of the form

$$\mathcal{N}(A) = \left\{ Q \ll P \big| f(Q|P) \le A \right\} \cap \mathcal{M}_e(S) =: \mathcal{Q}_A^f \cap \mathcal{M}_e(S)$$

for some f-divergence $f(Q|P) = E_P[f(Z_T)]$ with a convex function f and some constant A. (By the la Vallée-Poussin criterion, these sets are relatively weakly compact in $L^1(P)$ for $f = f^e$, or for $f = f^p$ with $0 < \gamma < \frac{1}{2}$.) Then we could use $\mathcal{Q}_A^f$ via (4.2) to generate an MCohUF Ψ_A^f which would produce, through the abstract approach, the same no-$\mathcal{A}$-good-deal valuation bounds as $\mathcal{N}(A)$ itself. However, we could also study $-\Psi_A^f$ as a risk measure on its own. The point is that Proposition 4.2 gives a way to generate possibly new and interesting examples of risk measures, whose generating set of measures naturally arises from no-good-deal valuation. This may deserve further exploration.

5 Conclusion and outlook

We present a method to define no-good-deal valuation bounds from an expected utility criterion, taking up ideas first presented in Hodges (1998) and further developed in Černý (2003). We provide a simple and rigorous approach to the static case for a general probability space, and show how to extend the basic idea rigorously to a dynamic situation. In a Lévy process setting, we demonstrate that constructing no-good-deal measures by imposing local conditional restrictions on their instantaneous market prices of risk yields no-good-deal valuation bounds with very good dynamic properties over time.

The dynamic results in Section 3 are given only for exponential utility, but all can be done equally well for logarithmic or power utility, $U^p(x) = \frac{x^{1-\gamma}}{1-\gamma}$ for $x > 0$ and $\gamma \in (0, 1)$. Without giving full details, we briefly sketch here how this works. First of all, one must extend the duality in Proposition 2.6 to the dynamic case like in Proposition 3.1; this is a straightforward direct calculation. The analogue of formula (3.3) for $f\left(Q^{(\beta,Y)}\big|P\right)$ is a bit more complicated, as explained below. But once one has overcome this, Proposition 3.4 and Theorem 3.7 extend in a straightforward manner. Finally, Theorem 3.9 mainly needs that the benchmark measure $\widehat{Q}$ is in $\mathcal{Q}^{\text{Lévy}}$, and so we want an extension of Esche and Schweizer (2005) to prove that the Lévy property of L is again preserved under the f-optimal martingale measure. This is provided by Theorem 2.6 of Jeanblanc, Klöppel, and Miyahara (2007) for U^p and Theorem 4.4.11 of Klöppel (2006) for U^ℓ, and existence of the f-optimal martingale measure for U^p is shown in Theorem 2.9 of Jeanblanc, Klöppel, and Miyahara (2007).

The only point not yet explained in this program is the generalisation of (3.3). As shown in Proposition 2.2 of Jeanblanc, Klöppel, and Miyahara (2007), $f\left(Q^{(\beta,Y)}\big|P\right)$ can like in (3.3) be written as an expectation of (exp of) the dt-integral over a deterministic function k^p applied to the Girsanov parameters β, Y. A small twist is that the expectation is no longer under $Q^{(\beta,Y)}$ but under a different measure $R(\beta, Y)$. However, all this does not matter since our local restrictions are again defined only in terms of the dt-integrands, and the latter are moreover deterministic in Theorem 3.9. Finally, the analogue of (3.3) for f^ℓ is given in Proposition 4.4.1 of Klöppel (2006).

A   Some key facts on Lévy processes

Based on Jacod and Shiryaev (1987), this section recalls some results on Lévy processes.

We work on a probability space $(\Omega, \mathcal{F}, P)$ with a finite time horizon T and a filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \le t \le T}$ satisfying the usual conditions. $\mathbf{P}$ is the predictable σ-field on $\Omega \times [0, T]$ and $\mathcal{B}^n$ the Borel σ-field on $\mathbb{R}^n$. For an $\mathbb{R}^n$-valued semimartingale X, μ^X denotes the random measure associated with its jumps and ν^P the predictable P-compensator of μ^X; only in this section, X is a process and not a payoff. We work throughout with a fixed but arbitrary truncation function $h : \mathbb{R}^n \to \mathbb{R}^n$ and denote by (B, C, ν) the P-characteristics of the semimartingale X with respect to h. We can and always do choose a version of the form

$$B = \int b \, dA, \qquad C = \int c \, dA, \qquad \nu(\omega; dx, dt) = K_{\omega,t}(dx) \, dA_t(\omega), \qquad (\text{A.1})$$

where A is a real-valued, predictable, increasing and locally integrable process; b is an $\mathbb{R}^n$-valued predictable process; c is a predictable process with values in the set of symmetric nonnegative definite $n \times n$ matrices; and $K_{\omega,t}(dx)$ is a transition kernel from $(\Omega \times [0, T], \mathbf{P})$ to $(\mathbb{R}^n, \mathcal{B}^n)$ such that $K_{\omega,t}(\{0\}) = 0$ and $\int_{\mathbb{R}^n} (1 \wedge |x|^2) K_{\omega,t}(dx) \le 1$.

Take $Q \in \mathcal{P}^{\ll}$ and an adapted $\mathbb{R}^n$-valued RCLL process $L = (L_t)_{0 \le t \le T}$ null at 0. L is a Q-*Lévy process* if for $s \le t$, the increment $L_t - L_s$ is Q-independent of $\mathcal{F}_s$ with a distribution depending only on $t - s$. For $Q = P$, we omit to mention P. Every Lévy process is a semimartingale, and a semimartingale L null at 0 is a $(P$-$)$Lévy process if and only if its $(P$-$)$characteristics are of the form

$$B_t = bt, \qquad C_t = ct, \qquad \nu^P(dx, dt) = K(dx) \, dt, \qquad (\text{A.2})$$

with $b \in \mathbb{R}^n$, c a symmetric nonnegative definite $n \times n$ matrix, and K a σ-finite measure on $(\mathbb{R}^n, \mathcal{B}^n)$ with $K(\{0\}) = 0$ and $\int_{\mathbb{R}^n} (1 \wedge |x|^2) K(dx) < \infty$. Briefly, the $(P$-$)$characteristics of L are then *deterministic and time-independent*.

We recall the next result to introduce the Girsanov parameters β, Y of $Q \in \mathcal{P}^{\ll}$.

Theorem A.1 (Jacod and Shiryaev (1987, Theorem III.3.24)) *Let X be a semimartingale with P-characteristics (B^P, C^P, ν^P) and corresponding processes c and A from (A.1). For any $Q \in \mathcal{P}^{\ll}$, there exist a $\mathbf{P} \otimes \mathcal{B}^n$-measurable function $Y \ge 0$ on $\Omega \times [0, T] \times \mathbb{R}^n$ and a predictable $\mathbb{R}^n$-valued process β satisfying*

$$\int_0^T \int_{\mathbb{R}^n} \left| (Y(s, x) - 1) h(x) \right| \nu^P(dx, ds) + \int_0^T |c_s \beta_s| \, dA_s + \int_0^T \beta_s^{\mathrm{tr}} c_s \beta_s \, dA_s < \infty \ Q\text{-a.s.}$$

and such that the Q-characteristics (B^Q, C^Q, ν^Q) of X are given by

$$B_t^Q = B_t^P + \int_0^t c_s \beta_s \, dA_s + \int_0^t \int_{\mathbb{R}^n} \left((Y(s, x) - 1) h(x) \nu^P(dx, ds) \right),$$

$$C_t^Q = C_t^P,$$

$$\nu^Q(dx, dt) = Y(t, x) \, \nu^P(dx, dt).$$

We call β and Y the Girsanov parameters *of Q (with respect to P, relative to X).*

Remark A.2 The Girsanov parameters are not unique; $Y(\omega, t, x)$ is unique only ν^P-a.e., and for fixed c and A we have A-a.e. uniqueness only for $c\beta$. In the sequel, we fix a P-Lévy process L and take all Girsanov parameters relative to L. We then identify all versions of Girsanov parameters (β, Y) which describe the same Q. In particular, we call Girsanov parameters *time-independent* if there exists a version with this property.

The following result gives a nice parametrisation for $\mathcal{P}^{\approx}$; it describes the density process of any $Q \in \mathcal{P}^{\approx}$ in terms of its Girsanov parameters (β, Y) and the Lévy process L.

Proposition A.3 (Esche and Schweizer (2005, Proposition 3)) *Suppose $\mathbb{F} = \mathbb{F}^L = (\mathcal{F}_t^L)_{0 \le t \le T}$ is the P-augmentation of the filtration generated by an $\mathbb{R}^n$-valued P-Lévy process L with characteristic triplet (b, c, K) as in (A.2). If $Q \in \mathcal{P}^{\approx}$ has Girsanov parameters (β, Y), the density process of Q with respect to P is given by $Z = \mathcal{E}(N^Q)$ with*

$$N_t^Q := \int_0^t \beta_s^{\mathrm{tr}} \, dL_s^c + \int_0^t \int_{\mathbb{R}^n} \big(Y(s, x) - 1\big)\big(\mu^L(dx, ds) - K(dx) \, ds\big) \text{ for } t \in [0, T], \quad \text{(A.3)}$$

where L^c is the continuous local martingale part of L and $\mathcal{E}(\cdot)$ the stochastic exponential.

References

[1] F. Bellini and M. Frittelli (2002), On the existence of minimax martingale measures, *Mathematical Finance* 12, 1–21

[2] A. E. Bernardo and O. Ledoit (2000), Gain, loss, and asset pricing, *Journal of Political Economy* 108, 144–172

[3] S. Biagini and M. Frittelli (2008), A unified framework for utility maximization problems: an Orlicz space approach, *Annals of Applied Probability* 18(3), 1–41

[4] T. Björk and I. Slinko (2006), Towards a general theory of good-deal bounds, *Review of Finance* 10, 221–260

[5] O. Bondarenko and I. R. Longarela (2004), Benchmark good-deal bounds: An application to stochastic volatility models of option pricing, preprint, Stockholm School of Economics, April 2004, `http://www2.hhs.se/personal/Rodriguez/research.htm`

[6]   P. Carr, H. Geman and D. B. Madan (2001), Pricing and hedging in incomplete markets, *Journal of Financial Economics* 62, 131–167

[7]   A. Černý (2003), Generalised Sharpe ratios and asset pricing in incomplete markets, *European Finance Review* 7, 191–233

[8]   A. Černý and S. Hodges (2002), The theory of good-deal pricing in financial markets, in: H. Geman, D. Madan, S. R. Pliska and T. Vorst (eds.), *Mathematical Finance – Bachelier Congress 2000*, Springer, 175–202

[9]   P. Cheridito and M. Kupper (2006), Time-consistency of indifference prices and monetary utility functions, preprint, Princeton University, January 2006, `http://www.math.ethz.ch/~kupper`

[10] A. S. Cherny (2007a), Pricing with coherent risk, preprint, Moscow State University, to appear in *Theory of Probability and Its Applications*,

[11] A. S. Cherny (2007b), Equilibrium with coherent risk, preprint, Moscow State University, to appear in *Theory of Probability and Its Applications*

[12] J. H. Cochrane and J. Saá-Requejo (2000), Beyond arbitrage: good-deal asset price bounds in incomplete markets, *Journal of Political Economy* 108, 79–119

[13] F. Delbaen (2006), The structure of m-stable sets and in particular of the set of risk neutral measures, *Séminaire de Probabilités XXXIX, Lecture Notes in Mathematics* 1874, Springer, 215–258

[14] F. Esche and M. Schweizer (2005), Minimal entropy preserves the Lévy property: How and why, *Stochastic Processes and Their Applications* 115, 299–327

[15] H. Föllmer and A. Schied (2004), *Stochastic Finance. An Introduction in Discrete Time*, 2nd edition, de Gruyter

[16] M. Frittelli (2000a), Introduction to a theory of value coherent with the no-arbitrage principle, *Finance and Stochastics* 4, 275–297

[17] M. Frittelli (2000b), The minimal entropy martingale measure and the valuation problem in incomplete markets, *Mathematical Finance* 10, 39–52

[18] T. Fujiwara and Y. Miyahara (2003), The minimal entropy martingale measures for geometric Lévy processes, *Finance and Stochastics* 7, 509–531

[19] T. Goll and L. Rüschendorf (2001), Minimax and minimal distance martingale measures and their relationship to portfolio optimization, *Finance and Stochastics* 5, 557–581

[20] S. Hodges (1998), A generalization of the Sharpe ratio and its applications to valuation bounds and risk measures, *FORC Preprint* 98/88, University of Warwick

[21] J. Jacod and A. N. Shiryaev (1987), *Limit Theorems for Stochastic Processes*, Springer

[22] S. Jaschke and U. Küchler (2001), Coherent risk measures and good-deal bounds, *Finance and Stochastics* 5, 181–200

[23] M. Jeanblanc, S. Klöppel and Y. Miyahara (2007), Minimal f^q-martingale measures for exponential Lévy processes, *Annals of Applied Probability* 17, 1615–1638

[24] Yu. M. Kabanov and C. Stricker (2002), On the optimal portfolio for the exponential utility maximization: remarks to the six-author paper, *Mathematical Finance* 12, 125–134

[25] C. Kardaras (2007), No-free-lunch equivalences for exponential Lévy models under convex constraints, preprint, Boston University, January 2007, to appear in *Mathematical Finance*

[26] S. Klöppel (2006), Dynamic valuations in incomplete markets, PhD thesis No. 16666, ETH Zürich, `http://www.math.ethz.ch/ kloeppel`

[27] S. Klöppel and M. Schweizer (2005), Dynamic utility indifference valuation via convex risk measures, NCCR FINRISK working paper No. 209, ETH Zürich, August 2005, `http://www.nccr-finrisk.unizh.ch/ wps.php?action=query&id=209`

[28] S. Klöppel and M. Schweizer (2007), Dynamic indifference valuation via convex risk measures, *Mathematical Finance* 17, 599–627

[29] D. Kramkov and W. Schachermayer (1999), The asymptotic elasticity of utility functions and optimal investment in incomplete markets, *Annals of Applied Probability* 9, 904–950

[30] F. Liese and I. Vajda (1987), *Convex Statistical Distances*, Teubner

[31] I. R. Longarela (2001), An extension of good-deal asset price bounds, SSE/EFI Working Paper No. 448, Stockholm School of Economics, `http://swopec.hhs.se/hastef/abs/hastef0448.htm`

[32] S. A. Ross (2005), *Neoclassical Finance*, Princeton University Press

[33] W. Schachermayer (2001), Optimal investment in incomplete markets when wealth may become negative, *Annals of Applied Probability* 11, 694–734

[34] J. Staum (2004), Fundamental theorems of asset pricing for good deal bounds, *Mathematical Finance* 14, 141–161

Susanne Klöppel
TU Wien
Finanz- und Versicherungsmathematik
Wiedner Hauptstrasse 8, 105-1 FAM
1040 Wien
Austria
kloeppel@fam.tuwien.ac.at

Martin Schweizer
ETH Zürich
Departement Mathematik
ETH-Zentrum, HG G 51.2
8092 Zürich
Switzerland
martin.schweizer@math.ethz.ch