Utility maximization, risk aversion, and stochastic dominance

Author(s):
Beiglböck, Mathias; Muhle-Karbe, Johannes; Temme, Johannes

Utility maximization, risk aversion, and stochastic dominance

Mathias Beiglböck · Johannes Muhle-Karbe · Johannes Temme

Abstract Consider an investor trading dynamically to maximize expected utility from terminal wealth. Our aim is to study the dependence between her risk aversion and the distribution of the optimal terminal payoff. Economic intuition suggests that high risk aversion leads to a rather concentrated distribution, whereas lower risk aversion results in a higher average payoff at the expense of a more widespread distribution. Dybvig and Wang (J. Econ. Theory, 2011, to appear) find that this idea can indeed be turned into a rigorous mathematical statement in one-period models. More specifically, they show that lower risk aversion leads to a payoff which is larger in terms of second order stochastic dominance. In the present study, we extend their results to (weakly) complete continuous-time models. We also complement an ad-hoc counterexample of Dybvig and Wang, by showing that these results are "fragile", in the sense that they fail in essentially *any* model, if the latter is perturbed on a set of arbitrarily small probability. On the other hand, we establish that they hold for power investors in models with (conditionally) independent increments.

1 Introduction

A classical problem in mathematical finance and financial economics is to maximize expected utility from terminal wealth. This means that—given a time horizon T and a utility function

M. Beiglböck · J. Temme
Fakultät für Mathematik, Universität Wien, Nordbergstrasse 15, 1090 Wien, Austria
e-mail: mathias.beiglboeck@univie.ac.at

J. Temme
e-mail: johannes.temme@univie.ac.at

J. Muhle-Karbe (✉)
Departement Mathematik, ETH Zürich, Rämistrasse 101, 8092 Zürich, Switzerland
e-mail: johannes.muhle-karbe@math.ethz.ch

U describing the investor's preferences—one tries to choose a trading strategy such that the terminal value $\hat{X}_T$ of the corresponding wealth process maximizes $\mathbb{E}\left[U(X_T)\right]$ over all wealth processes of competing strategies. Existence and uniqueness of the maximizer $\hat{X}$ are assured in very general models and under essentially minimal assumptions (cf., e.g., [10] and the references therein). However, much less is known about the qualitative properties of $\hat{X}$ and, in particular, about their dependence on the investor's attitude towards risk measured, e.g., in terms of her *absolute risk aversion* $-U''/U'$.

Since comparative statics for the composition of the investor's portfolio are impossible to obtain in any generality, Dybvig and Wang ([4], henceforth DW) have recently proposed to compare the *distributions* of the optimal payoffs instead. They show that—in one-period models—the payoffs of investors with ordered absolute risk aversion can be ranked in terms of stochastic dominance relationships. More specifically, suppose investor L is less risk averse than the more risk inverse investor M, and the corresponding optimal payoffs $\hat{X}_T^M$, $\hat{X}_T^L$ have finite first moments. Then [4, Theorems 3 and 7] assert that $\hat{X}_T^L$ dominates $\hat{X}_T^M$ in the *monotone convex order*,

$$\hat{X}_T^M \leq_{MC} \hat{X}_T^L, \tag{1.1}$$

that is, $\mathbb{E}[c(\hat{X}_T^M)] \leq \mathbb{E}[c(\hat{X}_T^L)]$ for every monotone increasing convex function $c : \mathbb{R}_+ \to \mathbb{R}$. Moreover, if either of the utility functions has nonincreasing absolute risk aversion, then DW also obtain the sharper assertion that

$$\mathbb{E}[\hat{X}_T^M] \leq \mathbb{E}[\hat{X}_T^L] \quad \text{and} \tag{1.2}$$

$$(\hat{X}_T^M - \mathbb{E}[\hat{X}_T^M]) \leq_C (\hat{X}_T^L - \mathbb{E}[\hat{X}_T^L]), \tag{1.3}$$

where $\leq_C$ denotes the *convex order*, i.e., (1.3) asserts that, for every convex function c, $\mathbb{E}[c(\hat{X}_T^M - \mathbb{E}[\hat{X}_T^M])] \leq \mathbb{E}[c(\hat{X}_T^L - \mathbb{E}[\hat{X}_T^L])]$. Both $\leq_C$ and $\leq_{MC}$ are *second order stochastic dominance* relations.

By Strassen's characterization of the convex order, cf. [17], (1.3) is tantamount to the existence of a random variable ε with $\mathbb{E}[\varepsilon|\hat{X}_T^M] = 0$ such that, in distribution,

$$\hat{X}_T^L = \hat{X}_T^M + (\mathbb{E}[\hat{X}_T^L] - \mathbb{E}[\hat{X}_T^M]) + \varepsilon. \tag{1.4}$$

In plain English, this means that the less risk averse investor is willing to accept the extra noise ε in exchange for the additional risk premium $\mathbb{E}[\hat{X}_T^L] - \mathbb{E}[\hat{X}_T^M] \geq 0$.

In addition, DW also construct some counterexamples showing that the above results generally do not hold in incomplete markets.

The purpose of the present study is threefold. Firstly, we prove an analogue of the main result of DW—which is stated in the discrete one-period setting common in much of economics—in the continuous-time framework prevalent in mathematical finance, under the assumption that the market is (weakly) complete. Whereas it would also be possible to extend the approach of DW, we believe that our presentation is both more compact and more transparent.

Next, in Sect. 3, we shed more light on the fragility of this structural result in incomplete markets. Whether the counterexamples of DW use somewhat ad-hoc models and utility functions, we show that—even for investors with power utilities—the result does not hold in *any* finite state model, if the latter is perturbed by adding just a single extra branch with arbitrarily small probability.

Finally, in Sect. 4, we take a look at additional structural assumptions which ensure the validity of DW's result also in incomplete markets. More specifically, we show that it holds

for power utility investors if the increments of the return processes are independent or, more generally, independent conditional on some stochastic factor process. We emphasize that even though the utility maximization problem can be solved fairly explicitly in this setup, the stochastic dominance relationship apparently cannot be read off the formulas. Instead, we prove the result by induction in a discrete approximation of the model and then pass to the limit. An extension of this result to more general preferences and/or market models appears to be a challenging topic for future research.

2 (Weakly) complete markets

Fix a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0,T]}, \mathbb{P})$. We consider a market of one riskless and d risky assets and work in discounted terms. That is, the riskless asset is supposed to be normalized to 1, whereas the (discounted) price process of the risky asset is assumed to be modeled by an $\mathbb{R}^d$-valued semimartingale S.

2.1 Utilities defined on the positive halfline

The investor's preferences are described by a *utility function*. Here we first consider the case where the latter is defined on the positive halfline. That is, it is assumed to be a strictly increasing, strictly concave, twice differentiable mapping $U : \mathbb{R}_+ \to \mathbb{R} \cup \{\infty\}$ satisfying the Inada conditions $\lim_{x \to \infty} U'(x) = 0$ and $\lim_{x \to 0} U'(x) = \infty$. Given a utility function U, the quotient $-U''(x)/U'(x)$ is called the *absolute risk aversion* of U at $x \in (0, \infty)$, cf. [1,13]. An investor with utility function U_M is called *more risk averse* than an agent with utility function U_L, written $U_M \ll U_L$, if the absolute risk aversion of U_M dominates the absolute risk aversion of U_L pointwise. In the sequel we frequently use that $U_M \ll U_L$ if and only if $U'_L(x)/U'_M(x)$ is monotone increasing for all $x \in (0, \infty)$.

For the remainder of this section, we suppose that the market is *complete*, i.e., that the set of equivalent (local) martingale measures is a singleton $\mathbb{Q}$,[1] and consider the problem of maximizing expected utility, $\sup_{X_T} \mathbb{E}[U(X_T)]$. Here, X_T runs though the terminal values of all wealth processes X that can be generated by self-financing trading starting from an initial endowment $x > 0$, and satisfy the admissibility condition $X \geq 0$. Throughout, we suppose that the supremum is finite, as, e.g., for utility functions that are bounded from above. Then, it is well-known (cf., e.g., [10, Theorem 2.0]) that there is a unique optimal wealth process $\hat{X}$ related to the martingale measure $\mathbb{Q}$ via the first-order condition

$$U'(\hat{X}_T) = y \frac{d\mathbb{Q}}{d\mathbb{P}}. \tag{2.1}$$

Here, the Lagrange multiplier y is a constant given by the marginal indirect utility of the initial capital x (cf., e.g., [10] for more details).

For a more risk averse investor with utility function U_M and a less risk averse investor with utility function U_L, we are now able to state our first main result, the counterpart of [4, Theorem 3] in continuous time: Lower risk aversion leads to a terminal payoff that is larger in the monotone convex order (1.1).

---

[1] In fact, an inspection of the proofs shows that it is sufficient to assume that the dual minimizer of [10] is the same for both agents. If this holds for *all* agents, this property has been called *weak completeness* of the financial market, see [11,16].

Theorem 2.1 *Let $\hat{X}_T^M$, $\hat{X}_T^L$ be integrable and suppose that $U_M \ll U_L$. Then*

$$\hat{X}_T^M \leq_{MC} \hat{X}_T^L.$$

If the absolute risk aversion of at least one agent is *nonincreasing*[2] we also obtain the stronger convex order result (1.3).

Theorem 2.2 *Let $\hat{X}_T^M$, $\hat{X}_T^L$ be integrable and suppose that $U_M \ll U_L$. If, in addition, either U_M or U_L has nonincreasing absolute risk aversion, then*

$$(\hat{X}_T^M - \mathbb{E}[\hat{X}_T^M]) \leq_C (\hat{X}_T^L - \mathbb{E}[\hat{X}_T^L]).$$

To simplify what has to be proved we use the following well-known characterization of the (monotone) convex order, which is a straightforward consequence of the monotone convergence theorem.

Lemma 2.3 *Let X, Y be random variables with finite first moments.*

  (i) *We have $X \leq_{MC} Y$ if and only if $\mathbb{E}\left[(X - K)^+\right] \leq \mathbb{E}\left[(Y - K)^+\right]$ for all $K \in \mathbb{R}$.*
  (ii) *We have $X \leq_C Y$ if and only if $X \leq_{MC} Y$ and $\mathbb{E}[X] = \mathbb{E}[Y]$.*

Proof of Theorem 2.1 To simplify notation, first notice that we may assume $y_M = y_L = 1$. Indeed this is achieved by rescaling U_M and U_L by the factor y_M and y_L, respectively, which has no affect on the utility maximization problem and the risk aversion of the utility functions. Setting $D := d\mathbb{Q}/d\mathbb{P}$, $F := (U_M')^{-1}$, and $G := (U_L')^{-1}$, the first-order condition (2.1) can be rewritten as

$$\hat{X}_T^M = F(D), \quad \hat{X}_T^L = G(D) \tag{2.2}$$

and $U_M \ll U_L$ implies that $F(x)/G(x)$ is decreasing in x.

Next, notice that there exists $q \in \mathbb{R}$ such that, almost surely,

$$q \leq \hat{X}^M \leq \hat{X}^L \quad \text{or} \quad q \geq \hat{X}^M \geq \hat{X}^L. \tag{2.3}$$

To see this consider $\rho := D(\mathbb{Q})$. Since the value processes are $\mathbb{Q}$-martingales by [10, Theorem 2.0] and have the same initial value x, it follows that

$$\int F \, d\rho = \int F(D) \, d\mathbb{Q} = \int \hat{X}^M \, d\mathbb{Q} = \int \hat{X}^L \, d\mathbb{Q} = \int G(D) \, d\mathbb{Q} = \int G \, d\rho.$$

As F and G are continuous this implies that there exists $p > 0$ such that $F(p) = G(p) =: q$. Since F/G is decreasing, we obtain for all $x \in \mathbb{R}$ that either $q \leq F(x) \leq G(x)$ or $q \geq F(x) \geq G(x)$, which yields (2.3).

As U_M' and U_L' are decreasing, we deduce from (2.3) and $U_M'(\hat{X}^M) = D = U_L'(\hat{X}^L)$ that $(D - p)(\hat{X}^L - \hat{X}^M) \leq 0$. The identity

$$0 = \mathbb{E}_{\mathbb{Q}}\left[\hat{X}^L - \hat{X}^M\right] = p\mathbb{E}\left[\hat{X}^L - \hat{X}^M\right] + \mathbb{E}\left[(D - p)(\hat{X}^L - \hat{X}^M)\right]$$

now yields the intermediate result $\mathbb{E}[\hat{X}^M] \leq \mathbb{E}[\hat{X}^L]$.

It remains to establish that $\mathbb{E}[(\hat{X}^M - K)^+] \leq \mathbb{E}[(\hat{X}^L - K)^+]$ for $K \in \mathbb{R}$. If $K \geq q$ this is a trivial consequence of (2.3).

If $K \leq q$, then (2.3) implies $\mathbb{E}[(\hat{X}^M - K)^-] < \mathbb{E}[(\hat{X}^L - K)^-]$. Adding the inequality $\mathbb{E}[\hat{X}^M - K] \leq \mathbb{E}[\hat{X}^L - K]$ we obtain the desired relation $\mathbb{E}[(\hat{X}^M - K)^+] \leq \mathbb{E}[(\hat{X}^L - K)^+]$ also in this case. □

---

[2] E.g., this holds for investors with power utility functions $x^{1-p}/(1 - p)$, i.e, with constant relative risk aversion $0 < p \neq 1$.

The proof of Theorem 2.2 follows a similar scheme.

Proof of Theorem 2.2 By Lemma 2.3 it suffices to show $\hat{X}^M \leq_{MC} \hat{X}^L - l$ where $l :=$ $\mathbb{E}[\hat{X}^L - \hat{X}^M]$. Note that $l > 0$ by the proof of Theorem 2.1.

Since either U_M or U_L has non increasing risk aversion, $U'_L(x + l)/U'_M(x)$ is increasing in x. As above we may assume $y_M = y_L = 1$. Setting $F := (U'_M)^{-1}$, $\tilde{G} := (U_L)^{-1} + l$ and $\tilde{\rho} := D(\mathbb{P})$ we have $\int F \, d\tilde{\rho} = \int \tilde{G} \, d\tilde{\rho}$. Arguing as before, we obtain the existence of a point $\tilde{q}$ such that, a.s.,

$$\tilde{q} \leq \hat{X}^M \leq \hat{X}^L - l \quad \text{or} \quad \tilde{q} \geq \hat{X}^M \geq \hat{X}^L - l$$

in analogy to (2.3).

As in the last step of the above proof of Theorem 2.1, this implies that $\mathbb{E}[(\hat{X}^M - K)^+] \leq \mathbb{E}[((\hat{X}^L - l) - K)^+]$ for all $K \in \mathbb{R}$. □

Remark 2.4 The converse of Theorem 2.1 also holds true: If two agents choose—in *every* complete market—payoffs $\hat{X}^M_T$, $\hat{X}^L_T$ satisfying

$$\hat{X}^M_T \leq_{MC} \hat{X}^L_T \tag{2.4}$$

then their corresponding utility functions satisfy $U_M \ll U_L$. This is a direct consequence of [4, Theorem 4], which establishes the above statement under the weaker assumption that (2.4) holds for all complete *one-period* market models.

2.2 Utilities defined on the entire real line

We now turn to investors with utility functions defined on the whole real line. Whereas the final results are analogous, the necessary definitions are technically more involved.

In this setting, we assume that the asset price process S is locally bounded. A *utility function* then is a strictly increasing, strictly concave, twice differentiable mapping $U : \mathbb{R} \to \mathbb{R} \cup \{\infty\}$ satisfying both the Inada conditions $\lim_{x \to \infty} U'(x) = 0$ and $\lim_{x \to -\infty} U'(x) = \infty$ and the condition of reasonable asymptotic elasticity:

$$\limsup_{x \to \infty} \frac{xU'(x)}{U(x)} < 1 \quad \text{and} \quad \liminf_{x \to -\infty} \frac{xU'(x)}{U(x)} > 1.$$

Following [15], the wealth process X of a self-financing trading strategy starting from an initial endowment $x \in \mathbb{R}$ is called *admissible*, if its utility $U(X_T)$ is integrable, and it is a supermartingale under all absolutely continuous local martingale measures $\mathbb{Q}$ with "finite V-expectation", $\mathbb{E}[V(d\mathbb{Q}/d\mathbb{P})] < \infty$, for the conjugate function $V(y) = \sup_{x \in \mathbb{R}}(U(x) - xy)$, $y > 0$, of U. Throughout, we suppose that the market admits an equivalent local martingale measure (i.e., satisfies NFLVR) and that for each $y > 0$, the dual problem $\inf_{\mathbb{Q}} \mathbb{E}[V(y d\mathbb{Q}/d\mathbb{P})]$ is finite with a dual minimizer $\hat{\mathbb{Q}}(y)$ in the set of equivalent local martingale measures. Sufficient conditions for the validity of the latter assumption can be found in [2]; in particular it holds if the market is complete or if the utility function under consideration is exponential, $U(x) = -e^{-\gamma x}$ with $\gamma > 0$, and an equivalent local martingale measure $\mathbb{Q}$ with finite entropy $\mathbb{E}[d\mathbb{Q}/d\mathbb{P} \log(d\mathbb{Q}/d\mathbb{P})] < \infty$ exists.

Subject to these assumptions, [15, Theorem 1] ensures that there is a unique wealth process $\hat{X}$ that maximizes utility from terminal wealth. Moreover, for a suitable Lagrange multiplier y, the latter is once again related to the corresponding dual minimizer $\hat{\mathbb{Q}}(y)$ via the first-order condition

$$y \frac{d\hat{\mathbb{Q}}(y)}{d\mathbb{P}} = U'(\hat{X}_T).$$

This was the key property for our proof of Theorems 2.1 and 2.2. Indeed, an inspection of the proofs shows that we never used that the domains of the utility functions are given by the positive halfline. Hence, we obtain the following analogous results:

Theorem 2.5 *Consider two agents with utility functions U_M, U_L defined on the whole real line and suppose the corresponding optimal terminal payoffs $\hat{X}_T^M$, $\hat{X}_T^L$ are integrable. Then if $U_M \ll U_L$ and the dual minimizers for both agents coincide, we have*

$$\hat{X}_T^M \leq_{MC} \hat{X}_T^L.$$

If, in addition, either U_M or U_L has nonincreasing absolute risk aversion, then

$$(\hat{X}_T^M - \mathbb{E}[\hat{X}_T^M]) \leq_C (\hat{X}_T^L - \mathbb{E}[\hat{X}_T^L]).$$

Since the so-called *minimal entropy martingale measure* is the dual minimizer for all exponential utility maximizers—irrespective of risk aversion and initial endowment—it follows that the above result is always applicable in this case. That is, no extra assumptions other than the integrability of the agents optimal payoffs need to be imposed on the financial market.

Corollary 2.6 *Consider two agents with exponential utilities $-e^{-\gamma_M x}$ resp. $-e^{-\gamma_L x}$. Then if $\gamma_M > \gamma_L$, the assumptions of both parts of Theorem 2.5 are always satisfied, provided that the agents optimal payoffs are integrable.*

Proof First notice that exponential utilities have constant and therefore nonincreasing absolute risk aversion. Next note that the notion of admissibility is both independent of the initial endowment and scale invariant for $U(x) = -e^{-\gamma x}$. Hence the optimal strategy is evidently independent of the initial endowment and inversely proportional to the absolute risk aversion γ. Since the Lagrange multiplier y is given by the marginal indirect utility (cf., e.g., [15, Theorem 1]), it then follows from the first-order condition that the dual minimizer is the same for all absolute risk aversions γ.   □

## 3 Structural counterexample in incomplete markets

In this section we show that—even in finite probability spaces and for investors with power utility functions—Theorems 2.1 and 2.2 are "fragile", in that they do not hold in *any* market model if the latter is perturbed by adding a single extra state with arbitrary small probability.

3.1 Basic idea

Our starting point is the simple observation that the monotone convex order between random variables X, Y can be destroyed by minimal perturbations of the distributions of X, Y, see Fig. 1.

To enforce that $\hat{X}_T^M \not\leq_{MC} \hat{X}_T^L$ it is sufficient[3] to find some number $K^* \in \mathbb{R}$ such that $\mathbb{E}[(\hat{X}_T^M - K^*)^+] > \mathbb{E}[(\hat{X}_T^M - K^*)^+]$. That is, in order to construct a model for which the montone convex order (1.1) fails, we want to assure that the more risk averse agent M attains with a higher probability than L large values above a certain threshold K^*. In particular, in finite probability spaces $\max_\omega \hat{X}_T^M(\omega) > \max_\omega \hat{X}_T^L(\omega)$ already assures that the montone convex order fails; this follows by considering the call with $K^* = \max_\omega \hat{X}_T^L$.

---

[3] Indeed, if $\hat{X}_T^M \not\leq_{MC} \hat{X}_T^L$ this is always witnessed by a hockeystick function, cf. Lemma 2.3.

(a) montone convex order   **(b)** no monotone convex order
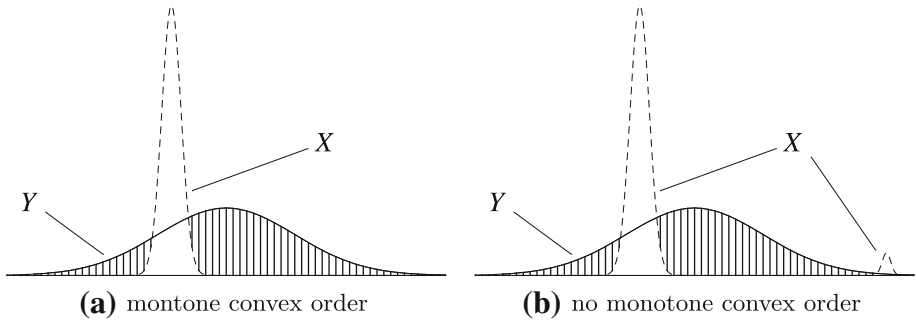
Fig. 1 Both illustrations show the distributions of random variables X and Y. **(a)** $X \leq_{MC} Y$; **(b)** $X \nleq_{MC} Y$

3.2 Concrete counterexample

We now present an explicit example for one riskless and one risky asset, showing that lower risk aversion does in general not lead to a larger portfolio in the monotone convex order. Moreover, our example exemplifies that this can happen even if the less risk averse investor always invest a larger fraction of her wealth in the risky asset. The latter is the decisive property used by DW for the proof of their results in incomplete one-period models.

We start with a complete two-period market model for which the monotone convex order (1.1) holds true by Theorem 2.1 above. We then alter this model by inserting a new branch after the first period which makes the model incomplete. The new branch occurs with an arbitrarily small probability ε so that the optimal strategies in the new model are almost identical to the original strategies. However, this new branch is constructed in such a way that the more risk averse agent M attains with positive probability a payoff that is larger than any possible payoff of L, which implies $\hat{X}_T^M \nleq_{MC} \hat{X}_T^L$.

For simplicity we consider a complete binomial model for which the stock does not change in the second period but stays constant.
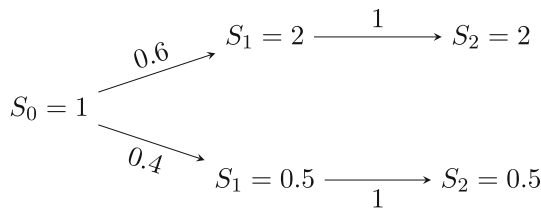

The preferences of agents M and L are given by power utilities with relative risk aversions $p_M = 0.9$ and $p_L = 0.3$, respectively.[4] By direct computation, we find that the agents optimally invest fractions $\hat{\pi}_0^M \approx 0.887$ resp. $\hat{\pi}_0^L \approx 1.853$ of their wealth in the risky asset at time $t = 0$. In particular, since L invests a larger fraction of wealth than M, agent L has less money than M when the price of the risky asset decreases.

The branch that we want to insert into S after the first period takes advantage of this disparity of wealth between M and L. It is given by S^*,
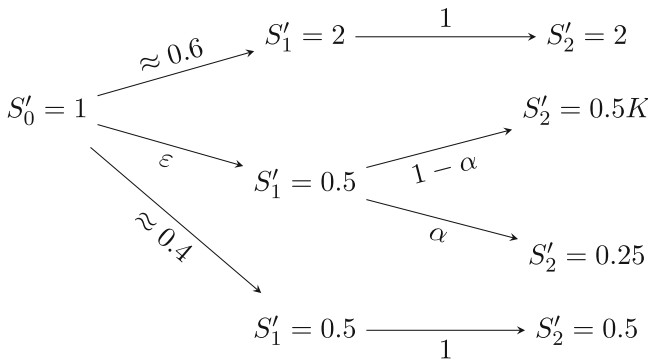
---

[4] Indeed any other choice of $p_i > 0$ can be made to work as well.

$$S_0^* = 0.5 \xrightarrow[\alpha]{1-\alpha} \begin{array}{c} S_1^* = 0.5K \\ \\ S_1^* = 0.25 \end{array}$$
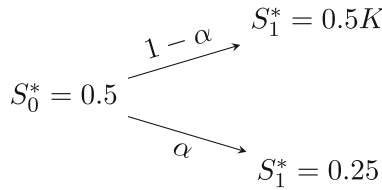
where we fix α to be a small probability, say $\alpha = 0.05$ and K a large stock value, for instance $K = 20$. Since this market offers both agents a very high probability of big fortune, M and L invest almost as much as admissibility allows and choose $\hat{\pi}^{M*} \approx 1.9492$, respectively $\hat{\pi}^{L*} \approx 1.9999$ as their optimal fraction of wealth invested in S^*.

We now "perturbate" S and define a new process S' by including S^* into S after the first step, i.e., we fix a small probability $\varepsilon = 0.01$ and define S' by

$$S_0' = 1 \begin{array}{c} \xrightarrow{\approx 0.6} S_1' = 2 \xrightarrow{1} S_2' = 2 \\ \\ \xrightarrow{\varepsilon} S_1' = 0.5 \xrightarrow[\alpha]{1-\alpha} \begin{array}{c} S_2' = 0.5K \\ S_2' = 0.25 \end{array} \\ \\ \xrightarrow{\approx 0.4} S_1' = 0.5 \xrightarrow{1} S_2' = 0.5 \end{array}$$

By the dynamic programing principle, the optimal trading strategies $\hat{\pi}_1'$ chosen at $t = 1$ for the new model S' are given by $\hat{\pi}^{M*}$, resp. $\hat{\pi}^{L*}$ above.[5] Since ε is small, the optimal trading strategies $\hat{\pi}_0'$ chosen at $t = 0$ are close to $\hat{\pi}_0^M$, resp. $\hat{\pi}_0^L$ above, and can be numerically computed to be given by $\hat{\pi}_1^{M'} \approx 0.8595$, resp. $\hat{\pi}_1^{L'} \approx 1.6622$.

We thus see that L invests in the second step a larger fraction of wealth in the stock than M. But since M's wealth after the first period is larger when $S_1' = 0.5$, M invests more *money* in S' than L. In particular we find that the optimal terminal payoffs $\hat{X}_2^M$, $\hat{X}_2^L$ satisfy $21.6897 \approx \max \hat{X}_2^M(\omega) > \max_\omega \hat{X}_2^L(\omega) \approx 6.5873$, where the maxima are attained at the event $S_2' = 0.5K$. Hence, $\hat{X}_2^M \not\preceq_{MC} \hat{X}_2^L$.

3.3 Counterexample in an n-period model

In the previous section we have seen in a concrete example that small changes of the model can cause the failure of the convex order relationship (1.1). Indeed this applies in a much wider setting; here we want to illustrate this in the case of an (arbitrage free) n-period model $(S_i)_{i=0}^n$ (where $n \geq 2$), defined on a finite probability space.

Consider, once again, agents M, L equipped with power utility functions with parameters $p_L = 0.3$ resp. $p_M = 0.9$. Assume for simplicity that the stock price does not stay constant

---

[5] In a general two period model the trading strategy chosen at $t = 1$ would depend on the current state of the first period. Since S_2' only changes in the branch given by S^* we can neglect this dependence.
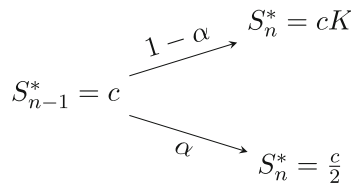
during any period. In this case the investors face a strictly concave optimization problem, hence, the optimal strategies $(\hat{\pi}_i^M)_{i=1}^n$, resp. $(\hat{\pi}_i^L)_{i=1}^n$ are uniquely determined. Denote by $\hat{X}^M$ resp. $\hat{X}^L$ the resulting optimal wealth processes.

We make the further assumption that the respective optimal wealth processes are not equal, more precisely that $\hat{X}_{n-1}^M$ is not equal to $\hat{X}_{n-1}^L$.

Fix an arbitrary small number $\eta > 0$. Then it is possible to replace the process S by a new process S' which agrees with S during the first $n-1$ stages and differs from S only in the last stage and with probability less than η, but for which (1.1) fails.

Using the assumption that $\hat{X}_{n-1}^M \neq \hat{X}_{n-1}^L$ we find that there exist $a, b, c \in \mathbb{R}, a > b$ such that the event $A = \{\hat{X}_{n-1}^M = a, \hat{X}_{n-1}^L = b, S_{n-1} = c\}$ has positive probability.

We now introduce a coin flip θ, independent of the stock price model and so that the outcome is $\{\theta = \text{head}\}$ with probability $\varepsilon < \eta$ and $\{\theta = \text{tail}\}$ with probability $1 - \varepsilon$. If the coin shows tail then the stock price process remains unchanged, i.e. $S' = S$. But in the event $A \cap \{\theta = \text{head}\}$, the stock price process in the last period is, as above, replaced by S^* given through

$$
S_{n-1}^* = c \quad
\begin{array}{c}
\overset{1-\alpha}{\nearrow} \quad S_n^* = cK \\[2em]
\underset{\alpha}{\searrow} \quad S_n^* = \frac{c}{2}
\end{array}
$$

where $K > 1$. An elementary analysis of the above example reveals that for α sufficiently close to 1 both agents will invest almost as much in the stock as admissibility allows. As $a > b$ we can arrange the constant K large enough so that the maximal payoff of agent M supersedes that of agent L as well as the maximum of $\hat{X}_n^L$.

There remains one issue to be dealt with: due to the change in the model we are now facing new optimal strategies and value processes. To cope with this problem we observe that the original problem was only changed on the portion $A \cap \{\theta = \text{head}\}$ of our space which has probability at most ε and that the possible gains on this set are bounded by some constant independent of ε. Consequently the perturbation of the original model vanishes as $\varepsilon \to 0$. As the original maximization problems had unique solutions, the new optimal strategies resemble the original ones as closely as we want (with the notable exception of the event $A \cap \{\theta = \text{head}\}$, in period $n-1$).

Summing up, upon choosing $\varepsilon > 0$ sufficiently small we obtain for the new optimal terminal wealth $\max_\omega \hat{X}_n^M(\omega) > \max_\omega \hat{X}_n^L(\omega)$ and in particular that (1.1) fails.

We conclude this section by pointing out that our argument is still valid if the stock is allowed to stay constant. Indeed, in this case the uniqueness of the optimal trading strategies is only violated in periods for which the stock does not change, but this does not affect the above reasoning.

4 Models with (conditionally) independent returns

In this last section, we consider some particular incomplete markets in continuous time, where the results of DW do hold.

More specifically, we focus on power utility investors in models of one riskless and one risky asset with independent or, more generally, conditionally independent returns. Its price process is first assumed to be modeled as the stochastic exponential $S = \mathcal{E}(R)$ of a *Lévy process R*, i.e, $dS_t/S_t = dR_t$. By its very definition, the Lévy process R has independent (and in fact also stationary) increments. Here, it can be interpreted as the returns process that generates the price process S of the risky asset in a multiplicative way.

Concerning preferences, we focus on investors with power utilities, i.e., $U(x) = x^{1-p}/(1-p)$, where $0 < p \neq 1$ denotes the investor's constant relative risk aversion. In this case, trading strategies are most conveniently parametrized in terms of the *fractions π_t* of wealth invested in the risky asset at time $t \in [0, T]$ (cf., e.g., [12] for a careful exposition of this matter). The wealth process corresponding to the risky fraction process $(\pi_t)_{t \in [0,T]}$ is then given by $dX_t/X_t = \pi_t dR_t$, i.e., $X_t = x\mathcal{E}(\int_0^{\cdot} \pi_s dR_s)_t$.

In this setting, it has been proved—by [14] in discrete time and, in increasing degree of generality, by [3,5,7,12] in continuous time—that the optimal policy is to invest a *constant* fraction $\hat{\pi}$ in the risky asset. The latter is known implicitly as the maximizer of some deterministic function, see [12]. In addition, it is possible to obtain some comparative statics for the optimal risky fractions here. More specifically, for two power utility functions $U_M \ll U_L$,[6] the optimal risky fractions $\hat{\pi}_M, \hat{\pi}_L$ satisfy $|\hat{\pi}_M| \leq |\hat{\pi}_L|$ and are non-negative (non-positive) if $\mathbb{E}[R_t]$ is non-negative (non-positive) for some (or equivalently all) t, cf. [18, Proposition 4.4].

The main result of this section is stated in the following theorem.

Theorem 4.1 *Suppose R is square-integrable and neither a.s. increasing nor a.s. decreasing, and $S = \mathcal{E}(R)$ is strictly positive. Then for power utility functions $U_M \ll U_L$ the optimal payoffs $\hat{X}^M$, $\hat{X}^L$ satisfy*

$$\hat{X}_T^M - \mathbb{E}[\hat{X}_T^M] \leq_C \hat{X}_T^L - \mathbb{E}[\hat{X}_T^M].$$

Since the optimal fractions are at least known implicitly as the maximizers of a scalar function and, in particular, satisfy $|\hat{\pi}_M| \leq |\hat{\pi}_L|$, one might think that this result can be obtained by a direct comparison of the corresponding wealth processes $\hat{X}_T^M = \mathcal{E}(\hat{\pi}_M R)_T$ and $\hat{X}_T^L = \mathcal{E}(\hat{\pi}_L R)_T$. However, the dependence of these random variables on the risky fractions is quite involved as can be seen by looking at the explicit formula for the stochastic exponential [6, Theorem I.4.61].

In a discrete-time setting, Theorem 4.1 can be established by induction using the results of Kihlstrom, Romer, and Williams [9] and the scaling properties of the power utilities. However, the corresponding result in continuous time cannot generally be obtained by passing to the limit since the continous-time optimizer can lead to bankruptcy if applied in discrete time, if it involves shortselling or leveraging the risky asset.

In order to prove Theorem 4.1, we therefore follow a different route. We first prove by induction the intermediate Proposition 4.2, which shows that the stochastic order holds true for discrete-time Euler approximations of $\hat{X}^M = \mathcal{E}(\hat{\pi}_M R)$ and $\hat{X}^L = \mathcal{E}(\hat{\pi}_L R)$. Theorem 4.1 is then established by showing that the stochastic dominance is preserved in the limit.

Proposition 4.2 *Let $(R_i)_i$ denote a sequence of i.i.d. random variables and let $\hat{\pi}_L, \hat{\pi}_M \in \mathbb{R}$ satisfying $\text{sgn}(\hat{\pi}_L) = \text{sgn}(\hat{\pi}_M)$ and $|\hat{\pi}_L| \geq |\hat{\pi}_M|$. Then*

$$\prod_{i=1}^{N} \left(1 + \hat{\pi}_M(R_i - \mathbb{E}[R_1])\right) \leq_C \prod_{i=1}^{N} \left(1 + \hat{\pi}_L(R_i - \mathbb{E}[R_1])\right) \quad \forall N \in \mathbb{N}.$$

---

[6] I.e., the relative risk aversion p^M of M is larger than its counterpart p^L for L.

Proof By induction on N. For $N = 0$ the assertion is trivial. For the induction step $N-1 \to N$ we apply Lemma 5.2 (using the independence of the R_i and the induction hypothesis) to obtain

$$\left(\prod_{i=1}^{N-1}(1 + \hat{\pi}_M(R_i - \mathbb{E}[R_1]))\right)(1 + \hat{\pi}_M(R_N - \mathbb{E}[R_1]))$$

$$\leq_C \left(\prod_{i=1}^{N-1}(1 + \hat{\pi}_L(R_i - \mathbb{E}[R_1]))\right)(1 + \hat{\pi}_M(R_N - \mathbb{E}[R_1])).$$

As Lemma 5.1 implies $(1 + \hat{\pi}_M(R_N - \mathbb{E}[R_1])) \leq_C (1 + \hat{\pi}_L(R_N - \mathbb{E}[R_1]))$, applying Lemma 5.2 once again proves the result. $\square$

Now we are in the position to prove Theorem 4.1:

Proof of Theorem 4.1 Since the optimal fraction for power utility is independent of the initial capital x, we set w.l.o.g. $x = 1$. By [18, Proposition 4.4], $|\hat{\pi}_M| \leq |\hat{\pi}_L|$ and $\hat{\pi}_M, \hat{\pi}_L$ are non-negative (non-positive) if $b := \mathbb{E}[R_1]$ is non-negative (non-positive). Thus, Proposition 4.2 implies

$$\prod_{i=1}^{N}\left(1 + \hat{\pi}_M\left(\Delta_i^N R - \frac{bT}{N}\right)\right) \leq_C \prod_{i=1}^{N}\left(1 + \hat{\pi}_L\left(\Delta_i^N R - \frac{bT}{N}\right)\right), \qquad (4.1)$$

where $\Delta_i^N R := R_{\frac{iT}{N}} - R_{\frac{(i-1)T}{N}}$ and $\mathbb{E}\left[\Delta_i^N R\right] = bT/N$. The left- and right-hand side of (4.1) are Euler approximations on an equidistant grid with mesh width T/N of the SDEs

$$d\bar{X}_t^i = \hat{\pi}_i \bar{X}_t^i \, d\bar{R}_t \qquad i = M, L,$$

where $\bar{R}_t = R_t - bt$. Since R is square-integrable, [18, Theorem A.3] shows that the Euler approximations converge in L^1 to the respective stochastic exponentials. As stochastic dominance is preserved under L^1-convergence we find

$$\mathcal{E}(\hat{\pi}_M \bar{R})_T \leq_C \mathcal{E}(\hat{\pi}_L \bar{R})_T.$$

Using $\mathcal{E}(\hat{\pi}_i \bar{R})_T = \mathcal{E}\left(\hat{\pi}_i R\right)_T \exp(-\hat{\pi}_i bT)$ and $\exp((\hat{\pi}_L - \hat{\pi}_M)bT) \geq 1$, Lemma 5.1 further implies

$$\mathcal{E}\left(\hat{\pi}_M R\right)_T \leq_C \mathcal{E}\left(\hat{\pi}_L R\right)_T - \left(\exp(\hat{\pi}_L bT) - \exp(\hat{\pi}_M bT)\right).$$

By $\mathbb{E}[\mathcal{E}\left(\hat{\pi}_i R\right)_T] = \exp(\hat{\pi}_i bT)$ the claim is proved. $\square$

4.1 Extension to models with conditionally independent increments

One can also extend Theorem 4.1 to somewhat more general models of the risky asset S.

Indeed, the proof of Theorem 4.1 and corresponding auxiliary results only use the independence of the increments of the Lévy process, but not their identical distribution. Hence, one can prove the convex order result of Theorem 4.1 along the same lines for processes $S_t = \mathcal{E}(R)_t$, where R has independent (but not necessarily identically distributed) increments. In this market model, the optimal policy for power utility is to invest a time-dependent but deterministic fraction $\hat{\pi}_t$ in the risky asset; the corresponding optimal wealth processes is then given by $\mathcal{E}(\int_0^{\cdot} \hat{\pi}_s dR_s)$.

This in turn allows to extend Theorem 4.1 also to models with *conditionally independent increments* (cf. [6, Chapter II.6] for more details). Loosely speaking, this means that

the return process R of the risky asset S is assumed to have independent increments with respect to an augmented filtration $\mathcal{G}_t$, that is, conditional on some stochastic factor processes. If these extra state variables are independent of the process driving the returns of the risky asset, Kallsen and Muhle-Karbe [8] show that the optimal policy is the same, both relative to the original and to the augmented filtration. With respect to the latter, one is dealing with a process with independent returns, such that Theorem 4.1 holds true. Statement (1.1) for the original filtration then follows immediately from the law of iterated expectations.

Appendix A: Auxiliary results on the convex order

In this Appendix, we state and prove two elementary results on the convex order, that are needed for the proof of Proposition 4.2.

Lemma 5.1 *Let X be a random variable and $a \geq 1$ a real number. Then*

$$X \leq_C aX - (a-1)\mathbb{E}[X].$$

Proof By Lemma 2.3, we have to prove

$$\mathbb{E}\left[(X-K)^+\right] \leq \mathbb{E}\left[(aX - (a-1)\mathbb{E}[X] - K)^+\right]$$

for all $K \in \mathbb{R}$. By centering, it is easily seen that it suffices to show the result for random variables X with $\mathbb{E}[X] = 0$.

Let $K < 0$. Since $\mathbb{E}[X] = 0$ and $\int_{-\infty}^{K} x\, \mathbb{P}^X(dx) \leq 0$ imply $\int_{K}^{\infty} x\, \mathbb{P}^X(dx) \geq 0$, we find

$$\mathbb{E}[(X-K)^+] \leq \int_{K}^{\infty} (x-K)\, \mathbb{P}^X(dx) + (a-1)\int_{K}^{\infty} x\, \mathbb{P}^X(dx),$$

$$\leq \int_{K/a}^{\infty} (ax-K)\, \mathbb{P}^X(dx) = \mathbb{E}\left[(aX-K)^+\right].$$

Next assume $K \geq 0$. As $\int_{K/a}^{K}(ax-K)\,\mathbb{P}^X(dx) \geq 0$, we conclude

$$\mathbb{E}\left[(X-K)^+\right] \leq \int_{K}^{\infty} (ax-K)\, \mathbb{P}^X(dx)$$

$$\leq \int_{\frac{K}{a}}^{\infty} (ax-K)\, \mathbb{P}^X(dx) = \mathbb{E}\left[(aX-K)^+\right].$$

$\square$

Lemma 5.2 *Let $X \leq_C Y$ and let Z be independent of X and Y. Then*

$$XZ \leq_C YZ.$$

Proof Let $c : \mathbb{R} \to \mathbb{R}$ be convex. Then the result easily follows from $\mathbb{P}^{ZX} = \mathbb{P}^Z \otimes \mathbb{P}^X$, $\mathbb{P}^{ZY} = \mathbb{P}^Z \otimes \mathbb{P}^Y$, and since the function $\tilde{c}(x) := c(zx)$ is again convex for all fixed $z \in \mathbb{R}$. $\square$

References

1. Arrow, K.J.: Aspects of the Theory of Risk-Bearing. Yrjö Jahnssonin Säätiö, Helsinki (1965)
2. Bellini, F., Fritelli, M.: On the existence of minimax martingale measures. Math. Financ. **12**(1), 1–21 (2002)
3. Benth, F.E., Karlsen, K.H., Reikvam, K.: Optimal portfolio selection with consumption and nonlinear integro-differential equations with gradient constraint: a viscosity solution approach. Financ. Stoch. **5**(3), 275–303 (2001)
4. Dybvig, P.H., Wang, Y.: Increases in risk aversion and the distribution of portfolio payoffs. J. Econ. Theory (2011) [To appear]
5. Framstad, N., Øksendal, B., Sulem, A.Y.: Optimal Consumption and Portfolio in a Jump Diffusion Market. Technical. Report 5, Norges Handelshøyskole. Institutt for foretaksøkonomi (1999)
6. Jacod, J., Shiryaev, A.N.: Limit Theorems for Stochastic Processes. 2nd edn. Springer, Berlin (2003)
7. Kallsen, J.: Optimal portfolios for exponential Lévy processes. Math. Meth. Oper. Res. **51**(3), 357–374 (2000)
8. Kallsen, J., Muhle-Karbe, J.: Utility maximization in models with conditionally independent increments. Ann. Appl. Probab. **20**(6), 2162–2177 (2010)
9. Kihlstrom, R., Romer, D., Williams, S.: Risk aversion with random initial wealth. Econometrica **49**(4), 911–920 (1981)
10. Kramkov, D., Schachermayer, W.: The asymptotic elasticity of utility functions and optimal investment in incomplete markets. Ann. Appl. Probab. **9**(3), 904–950 (1999)
11. Kramkov, D., Sirbu, M.: Sensitivity analysis of utility-based prices and risk-tolerance wealth processes. Ann. Appl. Probab. **16**(4), 2140–2194 (2006)
12. Nutz, M.: Power utility maximization in constrained exponential Lévy models. Math. Finance (2011) [To appear]
13. Pratt, J.W.: Risk aversion in the small and in the large. Econometrica **32**(1), 122–136 (1964)
14. Samuelson, P.A.: Lifetime portfolio selection by dynamic stochastic programming. Rev. Econ. Stat. **51**(3), 239–246 (1969)
15. Schachermayer, W.: A super-martingale property of the optimal portfolio process. Financ. Stoch. **7**(1), 433–456 (2003)
16. Schachermayer, W., Sîrbu, M., Taflin, E.: In which financial markets do mutual fund theorems hold true?. Financ. Stoch. **13**(1), 49–77 (2009)
17. Strassen, V.: The existence of probability measures with given marginals. Ann. Math. Stat. **36**(2), 423–439 (1965)
18. Temme, J.: Power utility maximization in discrete-time and continuous-time exponential Lévy models. Preprint (2011)