European Journal of Operational Research 242 (2015) 161–171

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Stochastics and Statistics

Risk-sensitive dividend problems

Nicole Bäuerle [a], Anna Jaśkiewicz [b,*]

[a] Department of Mathematics, Karlsruhe Institute of Technology, D-76128 Karlsruhe, Germany
[b] Wrocław University of Technology, PL-50-370 Wrocław, Poland

ARTICLE INFO

Article history:
Received 7 November 2013
Accepted 22 October 2014
Available online 3 November 2014

Keywords:
Markov decision process
Dividend payout
Risk aversion
History-dependent policy
Fixed point problem

ABSTRACT

We consider a discrete time version of the popular optimal dividend payout problem in risk theory. The novel aspect of our approach is that we allow for a risk averse insurer, i.e., instead of maximising the expected discounted dividends until ruin we maximise the expected *utility* of discounted dividends until ruin. This task has been proposed as an open problem in Gerber and Shiu (2004). The model in a continuous-time Brownian motion setting with the exponential utility function has been analysed in Grandits et al. (2007). Nevertheless, a complete solution has not been provided. In this work, instead we solve the problem in discrete time setup for the exponential and the power utility functions and give the structure of optimal history-dependent dividend policies. We make use of certain ideas studied earlier in Bäuerle and Rieder (2011), where Markov decision processes with general utility functions were treated. Our analysis, however, includes new aspects, since the reward functions in this case are not bounded.

1. Introduction

The dividend payout problem in risk theory has been introduced by de Finetti (1957) and has since then been investigated under various extensions during the decades up to now; see, for instance, Grandits, Hubalek, Schachermayer, and Žigo (2007) or Yao, Yang, and Wang (2011). The task is to find in a given model for the free surplus process of an insurance company, a dividend payout strategy that maximises the expected discounted dividends until ruin. Typical models for the surplus process are compound Poisson processes, diffusion processes, general renewal processes or discrete time processes. The reader is referred to Albrecher and Thonhauser (2009) and Avanzi (2009), where an excellent overview of recent results is provided.

In Gerber and Shiu (2004) the authors propose the problem of maximising the expected *utility* of discounted dividends until ruin instead of maximising the expected discounted dividends until ruin. This means that an insurance company is equipped with some utility function that helps it to measure the accumulated dividends paid to the shareholders. If this utility is increasing and concave, the company is risk averse (see Remark 2.2). To the best of our knowledge, there is only one work (Grandits et al., 2007), in which this idea was taken up. More precisely, Grandits et al. (2007) consider a linear Brownian motion model for the free surplus process and apply the exponential utility function to evaluate the discounted dividends until ruin. It turns out that the mathematics involved in the analysis of this problem is quite different from the one used in the risk neutral case and only partial results could be obtained. In contrast to the same problem with a risk neutral insurance company, where the optimal dividend payout strategy is of a barrier type (see e.g., Asmussen & Taksar, 1997), the authors in Grandits et al. (2007) are not able to identify the structure of the optimal dividend policy rigorously. They show imposing some further assumptions that there is a time dependent optimal barrier.

We study the same problem but with a discrete time surplus process. The risk neutral problem within such a framework can be found in Section 1.2 in Schmidli (2008) or in Section 9.2 in Bäuerle and Rieder (2011). By making use of the dynamic programming approach the authors in Bäuerle and Rieder (2011) and Schmidli (2008) prove that the optimal dividend payout policy is a stationary band-strategy. Albrecher, Bäuerle, and Thonhauser (2011), on the other hand, consider a discrete time model that is formulated with the aid of a general Lévy surplus process but the dividend payouts are allowed only at random discrete time points. This version can again be solved by the dynamic programming arguments. However, the problem with a general utility function is more demanding. Like in the continuous time setting (Grandits et al., 2007), it requires a sophisticated analysis. It is worth mentioning that Markov decision processes with general utility functions have been already studied in Kadota, Kurano, and Yasuda (1998) and Bäuerle and Rieder (2014). Moreover, there are also some papers, where the specific utility functions are considered. For example, Jaquette (1973, 1976) and Chung and Sobel (1987) are among the first who examined discounted payoffs in Markov decision processes with the decision maker that is equipped with a constant

* Corresponding author. Tel.: +48 71 320 3183.
E-mail addresses: nicole.baeuerle@kit.edu (N. Bäuerle), anna.jaskiewicz@pwr.edu.pl (A. Jaśkiewicz).

http://dx.doi.org/10.1016/j.ejor.2014.10.046

risk aversion, i.e., grades her random payoffs with the help of the exponential utility function. The common feature of all the aforementioned papers is the fact that they deal with bounded rewards or costs. Therefore, their results cannot be directly applied to our case, where the payoffs are unbounded. We make use of the special structure of the underlying problem and show that the optimal dividend payout policy is a time dependent band-strategy. The value function itself can be characterised as a solution to a certain optimality equation. Furthermore, we also study the dividend payout model with the power utility function. As noted in Bäuerle and Rieder (2014), the original Markov decision process can then be viewed as a Markov decision process defined on the extended state space. We employ these techniques to solve our model, but only in the first step, where we use an approximation of the value function in the infinite time horizon by value functions in the finite time horizons. In contrast to the exponential utility case, we can only partly identify the structure of the optimal dividend payout policy. However, we are able to show that there is a barrier such that when the surplus is above the barrier, it is always optimal to pay down to a state below the barrier. The value function is again characterised as a solution to some optimality equation. Summing up, the optimal dividend payout problem with the exponential utility function can be solved completely in the discrete time case, in contrast to the continuous-time problem in Grandits et al. (2007), whilst for the case with the power utility function we are at least able to identify the important global structure of the optimal policy.

The paper is organised as follows. In the next section we introduce the model together with mild assumptions and general history-dependent policies. Section 3 is devoted to a study of the exponential utility case. We show first that the value function J for discounted payoffs satisfies an optimality equation and give a lower and an upper bound for J. Then, we identify properties of the minimiser of the right-hand side of the optimality equation. This enables us to show that the minimiser indeed defines an optimal policy, which is a non-stationary band-policy. The non-stationarity is based only on the time-dependence. The power utility case is treated in Section 4. We pursue here a little different approach, but it also leads to an optimality equation. The policies obtained in this setting are really history-dependent. Nonetheless, we are still able to show that the optimal policy is of a barrier-type. In Section 5 we provide the policy improvement algorithm for the model with the exponential utility. Finally, Section 6 is devoted to concluding remarks and open issues.

2. The model

We consider the financial situation of an insurance company at discrete times, say $n \in \mathbb{N}_0 := 0, 1, 2, \ldots$. Assume there is an initial surplus $x_0 = x \in X := \mathbb{Z}$ and $x_0 \geq 0$. The surplus x_{n+1} at time $n+1$ evolves according to the following equation

$$x_{n+1} = x_n - a_n + Z_{n+1}, \;\text{ if } x_n \geq 0 \quad\text{and}\quad x_{n+1} = x_n, \;\text{ if } x_n < 0. \quad (2.1)$$

Here $a_n \in A(x_n) := \{0, \ldots, x_n\}$ denotes the dividends paid to the shareholders at time n, and Z_{n+1} represents the income (possibly negative) of the company during the time interval from n to $n+1$. More precisely, Z_{n+1} is the difference between premium and claim sizes in the $(n+1)$st time interval. Further, we assume that $Z_1, Z_2, \ldots$ are independent and identically distributed integer-valued random variables with distribution $(q_k)_{k\in\mathbb{Z}}$, i.e., $\mathbb{P}(Z_n = k) = q_k$, $k \in \mathbb{Z}$. A dividend payout problem in the risk theory can be viewed as a Markov decision process with the state space X, the set of actions $A(x)$ available in state x (for completeness, we put $A(x) = \{0\}$ for $x < 0$) and the transition probability $q(\cdot|x, a)$ of the next state, when x is the current state and a is the amount of dividend paid to the shareholders. Note that the dynamics of Eq. (2.1) implies that $q(y|x, a) = q_{y-x+a}$ for $x \geq 0$ and $q(x|x, a) = 1$ if $x < 0$. For the set of admissible pairs $D := \{(x, a) : \; x \in X, \; a \in A(x)\}$ we define the function $r : D \mapsto \mathbb{R}$ as $r(x, a) = a$ for $x \in X$.

The feasible history spaces are defined as follows $\Omega_0 = X$, $\Omega_k = D^k \times X$ and $\Omega_\infty = D^\infty$. A *policy* $\pi = (\pi_k)_{k\in\mathbb{N}_0}$ is a sequence of mappings from Ω_k to A such that $\pi_k(\omega_k) \in A(x_k)$, where $\omega_k = (x_0, a_0, \ldots, x_k) \in \Omega_k$. Let Γ be the class of all functions $g : X \mapsto A$ such that $g(x) \in A(x)$. A Markov policy is $\pi = (g_k)_{k\in\mathbb{N}_0}$ where each $g_k \in \Gamma$. By Π and Π^M we denote the set of all history-dependent and Markov policies, respectively. By the Ionescu–Tulcea theorem (Neveu, 1965), for each policy π and each initial state $x_0 = x$, a probability measure $\mathbb{P}_x^\pi$ and a stochastic process $(x_k, a_k)_{k\in\mathbb{N}_0}$ are defined on Ω_∞ in a canonical way, where x_k and a_k describe the state and the decision at stage k, respectively. By $\mathbb{E}_x^\pi$ we denote the expectation operator with respect to the probability measure $\mathbb{P}_x^\pi$.

Ruin occurs as soon as the surplus gets negative. The epoch τ of ruin is defined as the smallest integer n such that $x_n < 0$. The question arises as to how the risk-sensitive insurance company, equipped with some utility function will choose its dividend strategy. More precisely, we shall consider the following optimisation problem

$$\sup_{\pi\in\Pi} \mathbb{E}_x^\pi \, U_\gamma\left(\sum_{k=0}^{\infty} \beta^k r(x_k, a_k)\right) = \sup_{\pi\in\Pi} \mathbb{E}_x^\pi \, U_\gamma\left(\sum_{k=0}^{\tau-1} \beta^k a_k\right), \quad x \geq 0,$$

where $\beta \in (0, 1)$ is a discount factor and either

(1) U_γ is the exponential utility function, i.e., $U_\gamma(x) = \frac{1}{\gamma} e^{\gamma x}$ with $\gamma < 0$, or
(2) U_γ is the power utility function, i.e., $U_\gamma(x) = x^\gamma$ with $\gamma \in (0, 1)$.

Let Z be a random variable with the same distribution as Z_1. Throughout the paper the following assumptions will be supposed to hold true.

(A1) $\mathbb{E} Z^+ < +\infty$, where $Z^+ = \max\{Z, 0\}$;
(A2) $\mathbb{P}(Z < 0) > 0$.

Assumption (A2) allows to avoid a trivial case, when the ruin will never occur under any policy $\pi \in \Pi$.

Remark 2.1. In our study, we assume that the random variables $\{Z_n\}$ only take integer values and the initial capital is also integer. From the proof of Lemma 1.9 in Schmidli (2008), it follows that in our problem we can restrict without loss of generality to the integer dividend payments.

Remark 2.2. If the function U_γ is strictly concave and increasing as in our case, then the quantity $U_\gamma^{-1}(\mathbb{E}[U_\gamma(X)])$ is called a *certainty equivalent* of the random variable X. From the optimisation's point of view it does not matter which value $U_\gamma^{-1}(\mathbb{E}[U_\gamma(X)])$ or $\mathbb{E}[U_\gamma(X)]$ we study, because the inverse function U_γ^{-1} is monotonic. However, the certainty equivalent has an important meaning. If we apply the Taylor expansion, then the certainty equivalent can be written as follows

$$U_\gamma^{-1}(\mathbb{E}[U_\gamma(X)]) \approx \mathbb{E} X - \frac{1}{2} l(\mathbb{E} X) Var[X],$$

where

$$l(y) = -\frac{U_\gamma''(y)}{U_\gamma'(y)}$$

is called the *Arrow–Pratt* function of absolute risk aversion. Hence, the second term accounts for the variability of X (for a discussion see Bielecki & Pliska, 2003). If U_γ is concave like in our case, then $l(\cdot) \geq 0$ which means that the variance is subtracted. This fact implies that the decision maker is risk averse.

3. The exponential utility function

In this section we assume that the insurer is risk averse and grades her random payoffs by taking the expectations of the exponential utility function of these random rewards. More precisely, we assume that the decision maker is equipped with the constant risk coefficient