NRC Research Press

ARTICLE

Optimal forest management under financial risk aversion with discounted Markov decision process models

Mo Zhou and Joseph Buongiorno

Abstract: The common assumption of risk neutrality in forest decision making is generally inadequate because the stakeholders tend to be averse to fluctuations in the return criteria. In Markov decision processes (MDPs) of forest management, risk aversion and standard mean-variance analysis can be readily dealt with if the criteria are undiscounted expected values. However, with discounted criteria such as the fundamental net present value of financial returns, the classic mean-variance optimization is numerically intractable. In lieu of this, this paper (*i*) presents a linear-programming method to calculate the variance of discounted criteria conditional on any specific policy and (*ii*) adopts, as an alternative to the variance measure of risk, the "discount normalized variance" (DNV), an economically meaningful criterion consistent with income-smoothing behavior. The DNV is then used in procedures analogous to mean-variance analysis and certainty-equivalent optimization tractable by quadratic programming. The methods are applied to the management of uneven-aged, mixed-species forests in the southern United States. The results document the trade-off between the expected net present value and risk of financial returns, as well as the consequences for selected ecological criteria.

Key words: uneven-aged management, risk aversion, economics, multiple criteria, Markov decision process, quadratic programming.

Résumé : L'hypothèse répandue de la neutralité à l'égard du risque dans la prise de décision forestière est généralement inadéquate puisque les acteurs ont tendance à s'opposer aux fluctuations des critères de rendement. Dans les modèles de processus de décision de Markov (PDM) en aménagement forestier, les analyses d'aversion au risque et de variance moyenne standard peuvent être facilement effectuées si les critères sont des valeurs attendues non actualisées. Cependant, avec des critères actualisés, tels que la valeur actuelle nette fondamentale des rendements financiers, l'optimisation classique de la variance moyenne est numériquement insoluble. En revanche, cet article (*i*) présente premièrement une méthode de programmation linéaire pour calculer la variance des critères actualisés selon aucune politique spécifique et (*ii*) adopte, à titre d'alternative à la variance de la mesure du risque, la « variance normalisée escomptée » (VNE), un critère économiquement significatif compatible avec le comportement de lissage des revenus. La VNE est ensuite utilisée dans des procédures analogues à l'analyse de variance moyenne et à l'optimisation d'équivalent certain qui peuvent être résolues par la programmation quadratique. Les méthodes sont appliquées à l'aménagement des forêts mélangées inéquiennes dans le sud des États-Unis. Les résultats documentent le compromis entre la valeur actuelle nette et le risque des rendements financiers, ainsi que les conséquences pour certains critères écologiques. [Traduit par la Rédaction]

Mots-clés : aménagement inéquienne, aversion au risque, économie, critères multiples, processus de décision de Markov, programmation quadratique.

Introduction

The assumption of risk-neutral decision making is common in forest management, and the solution methods are well developed. It is nevertheless inappropriate when the stakeholders are sensitive to fluctuations in the return criteria and, under some circumstances, are willing to settle for a non-optimal expected return in exchange for reduced risk. This is, in particular, the case of nonindustrial private forest owners who tend to differ widely in terms of their financial risk tolerance (e.g., Lunnan et al. 2006; Andersson and Gong 2010; Tian et al. 2015; Feliciano et al. 2017; Laakkonen et al. 2018).

Recent works to address this issue in the context of optimal harvesting include Eyvindson and Kangas (2016), Couture et al. (2016), and Buongiorno et al. (2017), among others. Notwithstanding differences in risk measures, a consistent finding is that financial risk aversion leads to a cutting cycle (or rotation) shorter than the optimal one determined with the risk-neutral assumption. In the case of multistand forest planning, Tahvonen and Kallio (2006), Couture and Reynaud (2011), and Roessiger et al. (2011) conclude that natural or near-natural forests may arise from risk aversion due to diversification of age classes or structures.

One frequently used approach in dealing with decision making under risk is the family of Markov decision processes (MDPs) (Puterman 2009). MDPs with expected average or discounted criteria have often been applied in forest planning (e.g., Lembersky and Johnson 1975; Martell 1980; Buongiorno 2001; Zhou and Buongiorno 2011).

Received 13 December 2018. Accepted 14 April 2019.

M. Zhou.* Department of Forestry and Natural Resources, Purdue University, 715 W. State St., West Lafayette, IN 47906, USA.
J. Buongiorno. Department of Forest and Wildlife Ecology, University of Wisconsin-Madison, 1630 Linden Drive, Madison, WI 53705, USA.
Corresponding author: Mo Zhou (email: mozhou@purdue.edu).

*Mo Zhou currently serves as an Associate Editor; peer review and editorial decisions regarding this manuscript were handled by S. Ellen Macdonald and Bin Mei.

Published at www.nrcresearchpress.com/cjfr on 18 April 2019.

Risk-sensitive MDPs with expected average (undiscounted) criteria, first introduced in Howard and Matheson (1972) and Jaquette (1976), have been further analyzed by, for instance, Sobel (1994) and Fleming and Hernández-Hernández (1997, 1999). The much more difficult problem of MDPs with discounted criteria, investigated by Sobel (1982), Chung and Sobel (1987), and Levitt and Ben-Israel (2001), is still the subject of active research (e.g., Mannor and Tsitsiklis 2011; Guo et al. 2012; Xia 2018).

For applications to risk-averse conditions with undiscounted criteria, a quadratic programming (QP) formulation of MDPs is particularly attractive. It allows for constraints to reflect multiple objectives, and large problems can readily be solved with available quadratic programming algorithms (Filar et al. 1989). In forestry, Buongiorno et al. (2017) use quadratic programming models that, based on the variances of the undiscounted economic and ecological criteria, obtain optimal harvesting policies for risk-seeking and risk-averse landowners, using either mean-variance (Markowitz 1952) or certainty-equivalent (Freund 1956; Levitt and Ben-Israel 2001) theory.

Undiscounted criteria, however, fail to reflect the time preference of decision makers and the opportunity cost of investment alternatives, but discounting in mean-variance optimization of MDPs leads to hard nonlinear problems with no guarantee of convergence (Sobel 1982; Mannor and Tsitsiklis 2011), as the variance of MDPs with discounted criteria lacks the mathematical properties required for an optimal stationary policy (Denardo 1967; Sobel 1975, 1982). Thus, mean-variance or certainty-equivalent optimization is not feasible with discounted criteria such as the net present value (NPV) of financial returns commonly used in forest management.

The objective of this study was to seek ways of dealing with risk aversion and discounted criteria in forest management MDPs. First, based on Sobel (1982), we show how the variance of a discounted criterion resulting from any policy in an MDP framework can be calculated by linear programming. Second, we adopt Filar et al.'s (1989) "discount normalized variance" measure of risk, the variation in the difference between the periodical return and the long-run expected annual return. With this measure, a decision-making problem with risk aversion and discounted criteria can be expressed as a "variance-penalized optimization" (Filar et al. 1989), analogous to the certainty-equivalent optimization with undiscounted criteria. Following the description of these methods, the paper describes their application to the management of uneven-aged, mixed-species forests in the southern United States (US). The results show the trade-off between expected NPV and risk of financial returns, as well as the consequences of particular choices of financial risk – NPV combinations for undiscounted ecological criteria.

Methods

Markov decision process

In general, an MDP with discounted criteria is a 5-tuple (S, P, D, R, β), where S is a set of finite states $\{s_t\}$, P is a transition probability matrix containing entries of $p_{ij} = p(s_{t+1} = j|s_t = i)$, D is a set of decisions, R is a real-valued reward function, and β is the discount factor. A policy π is a mapping from state space S to decision space D.

In the context of this paper, S is a set of composite states of the forest stand condition and market level. Assuming that the two are independent, the composite transition probability is the product of the probability of stand growth between states and of market changes. A decision is a change from one stand state to another due to a harvest. For a decision, the monetary reward is determined by the amount of harvests and the market level. Ecological rewards are only associated with the postharvest forest condition.

The maximum expected sum of discounted values of monetary rewards over an infinite planning period, i.e., the expected NPV, E(NPV), of harvests, is obtained by solving the following linear programming problem (d'Epenoux 1963):

$$(1)\qquad \max_{y_{ik}} E(\mathrm{NPV}) = \sum_{i,k} R_{ik} y_{ik}$$

subject to

$$\sum_k y_{jk} - \beta \sum_{i,k} y_{ik} p(j|i,k) = \alpha_j \quad j = 1, \ldots, N$$

$$y_{ik} \geq 0 \quad i = 1, \ldots, N$$

in which the decision variable y_{ik} is the discounted time of being in stand-market state i and making decision k; R_{ik} is the instantaneous return associated with the state–decision pair (i, k), a function of stumpage price in state i and the volume harvested due to action k; β is the discount factor, $1/(1 + r)$, where r is the annual interest rate, assumed to be 3% in the case study; α_j is the initial probability of state j; $p(j|i, k)$ is the probability of the state transitioning to j in one year, given current state i and decision k. After solving eq. 1, the optimal policy is given by

$$(2)\qquad d_{ik} = \frac{y_{ik}}{\sum_k y_{ik}} \quad \forall i, k$$

where d_{ik} is the probability of decision k given state i. In the absence of constraint, $d_{ik} = 0$ or 1, the policy is deterministic and stationary: there is only one decision for each system state.

Variance of discounted criteria

The expected discounted value, v_i, of a criterion for any given policy $\{d_{ik}\}$ and initial state i is obtained by solving the following system of linear equations (Hillier and Lieberman 2005, p. 919):

$$(3)\qquad v_i = \sum_k d_{ik} R_{ik} + \beta \sum_j p(j|i, d_{ik}) v_j \quad \forall i$$

Given the expected discounted returns by the initial state, v_i, obtained by eq. 3, the variance, V_i, of the discounted value of returns conditional on a policy and initial state is then obtained by solving the following system of linear equations (Sobel 1982):

$$(4)\qquad V_i = \beta^2 \left\{ \sum_j p(j|i, d_{ik}) v_j^2 - \left[\sum_j p(j|i, d_{ik}) v_j \right]^2 \right\} + \beta^2 \sum_j p(j|i, d_{ik}) V_j \quad \forall i$$

Lastly, the variance of the NPV conditional on the probability distribution of the initial states, $\{\alpha_i\}$, is obtained with the general conditional variance formula (Lindgren 1968 p. 118):

$$(5)\qquad \mathrm{var}(X) = \mathrm{var}[E(X|Y)] + E[\mathrm{var}(X|Y)]$$

which in this application becomes

$$(6)\qquad \mathrm{var}(\mathrm{NPV}) = \mathrm{var}[E(\mathrm{NPV}|i)] + E[\mathrm{var}(\mathrm{NPV}|i)]$$

or

$$(7)\qquad \mathrm{var}(\mathrm{NPV}) = \mathrm{var}(v) + E(V)$$

with

$$(8)\qquad \mathrm{var}(v) = \sum_i \alpha_i v_i^2 - \left(\sum_i \alpha_i v_i\right)^2 \text{ and } E(V) = \sum_i \alpha_i V_i$$

Discount normalized variance penalized MDP

Attempting a mean-variance optimization such as maximizing the expected NPV with a constraint on the variance would require solving eqs. 3 to 8 simultaneously, but this is a highly nonlinear problem with no readily available algorithm and no guarantee of convergence. In view of this difficulty, Filar et al. (1989) proposed an alternative measure of variability, the discount normalized variance (DNV) associated with a policy π, defined as

$$(9)\qquad \mathrm{DNV}_\pi = \sum_{t=1}^{\infty} \beta^{t-1} E_\pi\{[R_t - (1-\beta)E(\mathrm{NPV}_\pi)]^2\}$$

in which E_π is the expected value operator with policy π and R_t is the random instantaneous reward at time t.

In eq. 9, $(1 - \beta)E(\mathrm{NPV}_\pi)$ is the approximate annualized NPV with policy π, because $(1 - \beta) \approx r$. Thus the DNV can simply be interpreted as the sum of discounted expected squared differences between actual return and expected annualized return over an infinite time horizon. In forestry, the expected annualized return from timber production can be viewed as a reference point for landowners to measure gain and loss, and reducing variations around it is consistent with income- smoothing behavior (Morduch 1995) as part of the risk-averse landowners' attempt to smooth consumption due to constraints on insurance and credit.

With this DNV measure of risk, Filar et al. (1989) proposed a "variance-penalized optimization" (VPO) that maximizes the weighted sum of the expected NPV and its discount penalized variance, analogous to the certainty-equivalent criterion for undiscounted criteria (Freund 1956; Levitt and Ben-Israel 2001):

$$(10)\qquad \mathrm{VPO} = \max_\pi[E(\mathrm{NPV}) - \lambda\mathrm{DNV}]$$

where λ is a nonnegative parameter that represents the risk-aversion level.

Filar et al. (1989) showed that the eq. 10 problem and the following variance-penalized MDP

$$(11)\qquad \mathrm{VPO} = \max_{y_{ik}}\left\{\sum_{i,k} R_{ik}y_{ik} - \lambda\left[\sum_{i,k} R_{ik}^2 y_{ik} - (1-\beta)\left(\sum_{i,k} R_{ik}y_{ik}\right)^2\right]\right\}$$

are related in that the optimal solution of eq. 11, y_{ik}^*, leads to the optimal policy, π^*, for eq. 10.

The eq. 11 model is a quadratic program, attractive for its ease of computation for large practical problems. Furthermore, it is analogous to the certainty-equivalence model with undiscounted objectives (Freund 1956; Levitt and Ben-Israel 2001) with the interpretation that

$$(12)\qquad \mathrm{DNV} = \sum_{i,k} R_{ik}^2 y_{ik} - (1-\beta)\left(\sum_{i,k} R_{ik}y_{ik}\right)^2$$

In the following application, the DNV (eq. 12) was used in objective functions as in eq. 11 and in constraints for mean-variance analysis. To compare the two measures of variability, var(NPV) and DNV, after a solution was obtained with the DNV criterion, the corresponding variance of the NPV was computed with eqs. 3 to 8.

Forest and market models

The methods outlined above were applied with an MDP of mixed loblolly pine and hardwood forests in the southern US. There were 64 forest states, defined by the basal area (low or high) in pines and hardwood trees of three size classes (pulpwood, small sawtimber, large sawtimber). The transition probabilities between stand states were those reported in Zhou and Buongiorno (2006). The market state was defined by a price index (low, medium, or high), based on the prices of softwoods and hardwoods from 1977 to 2014, and the transition probabilities between market states were obtained from Buongiorno et al. (2017). Combining the transition probabilities of stand states and market states gave the 192 × 192 transition probability matrix between all possible stand–market states.

Ecological criteria

While discounting is well established in economics and finance and the use of NPV is widely accepted, this is not generally the case for ecological criteria. There is still no consensus on whether the current and future environment and nature should have different weights and what should be the discount factor, if any (Arrow et al. 1996, 2014; Portney and Weyant 1999; Goulder and Stavins 2002; Hardisty and Weber 2009; Harrison 2010). Here, as in some previous studies (e.g., Buongiorno and Zhou 2015), each ecological criterion such as tree species diversity was measured with its undiscounted expected value over an infinite horizon, a value directly comparable with its current value. Consequently, only the long-term expected value of ecological criteria was considered in this study, in view of tracing the consequences of financial objectives and financial risk aversion on the long-term ecological condition of the forest.

In the MDP, a decision consisted of changing the stand state with a harvest or leaving it intact. Each postdecision stand state corresponded to two stand diversity criteria: tree species diversity and tree size diversity measured with Shannon's index (Zhou and Buongiorno 2006). Other indices of the stand condition were the postharvest basal area, the amount of CO_2e stored in the residual stand, the annual removals in sawtimber and pole timber, and the expected interval between harvests (cutting cycle), the inverse of the harvest frequency (Kaya and Buongiorno 1987). The long-term condition of the forest landscape (landscape diversity) was measured with a Shannon index based on the frequency of the various stand states in the steady state resulting from the continuous application of a particular harvesting policy (Zhou and Buongiorno 2006).

Results

Mean-variance analysis

Figure 1 shows, for the data used in this paper, the relationship between risk and the expected value of financial returns, using both the variance of the NPV, the classic measure of risk in finance, and the DNV suggested by Filar et al. (1989). Each pair of points on the graph was obtained by minimizing the DNV subject to a constraint on expected NPV and then computing the corresponding variance of the NPV with eqs. 3 to 9.

As seen in Fig. 1, both measures of financial risk increased at an increasing rate as the expected NPV increased from zero to its maximum unconstrained value of nearly US\$14 000·ha^{-1}. For a given value of expected NPV, the DNV was less than the variance of NPV; however, there was a strong linear relationship between the logarithm of the variance of NPV and the logarithm of the DNV, with a slope nearly equal to unity (Fig. 2). Thus, a relative change in var(NPV) corresponded to an approximately equal relative change in DNV, confirming that the DNV could be used as a proxy for variance in mean-variance analysis.

This is illustrated in Fig. 3 where a mean DNV frontier analogous to a mean-variance frontier was obtained by maximizing the expected NPV with the DNV less than or equal to a fraction of its level at maximum unconstrained NPV. Any point below the frontier was suboptimal because it provided less than the highest NPV for the level of risk represented by the DNV, and any point above

Fig. 1. Relation between the minimum discounted normalized variance, DNV, of financial returns and the corresponding variance of net present value, var(NPV), conditional on expected net present value, *E*(NPV). The data points at each level of expected net present value show the minimum discounted normalized variance of returns and the corresponding variance of their net present value. [Colour online.]

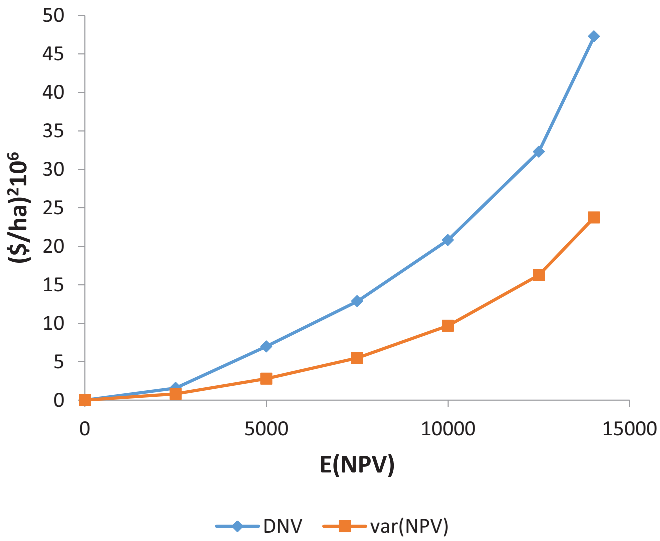

Fig. 2. Relation between the logarithm of the discount normalized variance of financial returns, ln(DNV), and the logarithm of the variance of net present value, ln(var(NPV)), with the data in Fig. 1. [Colour online.]

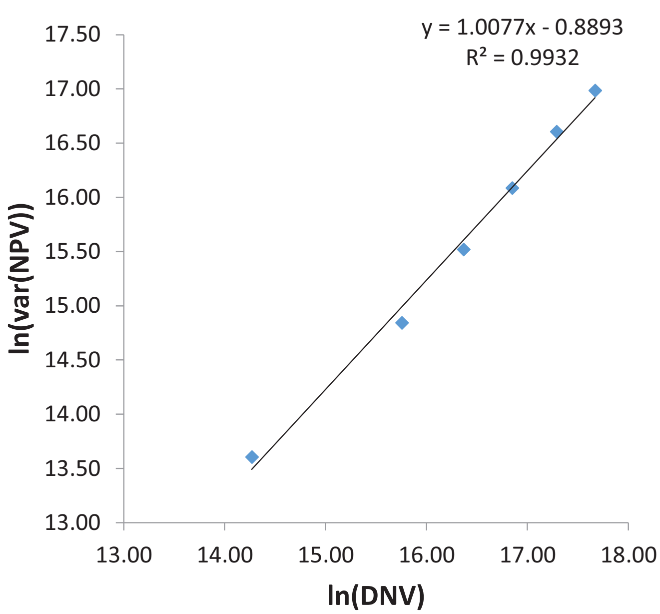

it was infeasible. With the data of this study, a 25% reduction in the DNV from its highest value resulted in a reduction of expected NPV by approximately US$1000·ha^{-1}, while a 75% reduction of the DNV caused the expected NPV to drop by half (Fig. 3). The extreme case of zero financial risk could only be achieved with no harvest and zero NPV.

The consequences of two choices along the risk-expectation frontier are shown in Table 1 for financial and ecological criteria. One choice assumed a risk-neutral decision maker who wanted to maximize the expected NPV of financial returns. The other choice was a risk-averse decision maker who wanted to halve the risk, expressed by the DNV, compared with what it was without risk aversion.

Fig. 3. Mean discount normalized variance frontier for financial return. Each point on the graph is the maximum expected net present value of returns, *E*(NPV), conditional on a level of the discount normalized variance of returns, DNV. The discounted normalized variance is relative to the DNV when the unconstrained expected NPV is maximized. [Colour online.]

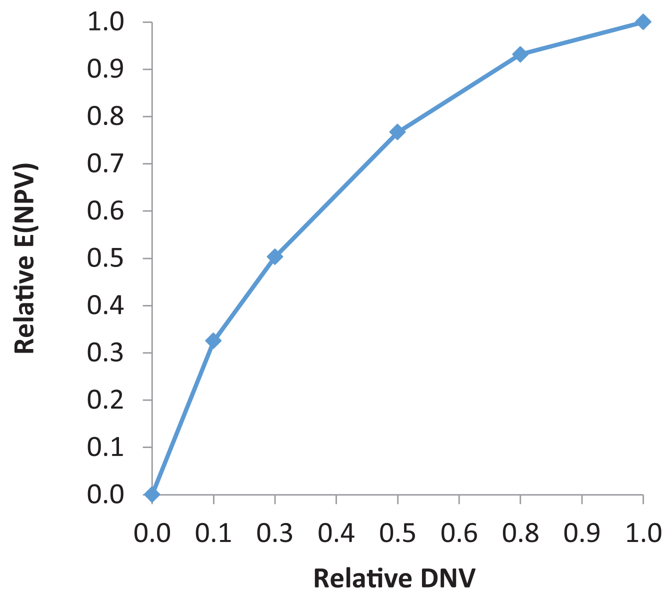

Table 1. Effects of financial risk aversion on management criteria.

Criteria	Risk neutral (1)	Risk averse (2)	Relative difference (2/1)
Discounted financial			
E(NPV) (\$·ha^{-1})	14024	10754	0.77
DNV (\$·ha^{-1})^2	47260078	23630039	0.50
var(NPV) (\$·ha^{-1})^2	23743580	10522139	0.44
Undiscounted financial			
Annual returns (\$·ha^{-1}·year^{-1})	321	255	0.79
Sawlog harvest (m^3·ha^{-1}·year^{-1})	5.2	4.9	0.94
Pole harvest (m^3·ha^{-1}·year^{-1})	2.4	2.5	1.07
Total harvest (m^3·ha^{-1}·year^{-1})	7.6	7.4	0.98
Undiscounted ecological			
Species diversity (index)	0.9	0.8	0.96
Size diversity (index)	1.8	1.8	1.00
Basal area (m^2·ha^{-1}·year^{-1})	12.8	12.1	0.94
CO_2e stock (t·ha^{-1})	103.6	98.9	0.95
Landscape diversity (index)	1.71	1.19	0.73
Cutting cycle (year)	5.5	5.2	0.95

Note: The risk-neutral solution maximized the expected net present value of financial returns, *E*(NPV), without constraint. The risk-averse solution constrained the discount normalized variance of the financial net present value, DNV, to half of its level in the risk-neutral case. Currency: US dollars.

Solution of the quadratic programming problem maximizing the expected NPV with or without the DNV constraint showed a reduction of expected NPV of 23% due to risk aversion and a corresponding reduction of expected undiscounted annual returns of 21%. The variance of the NPV obtained with eqs. 3 to 9 was reduced by 56%, similar for practical purposes to the DNV reduction of 50%.

The other large relative effect of risk aversion was the reduction of the landscape diversity index by 19%. The species diversity was reduced by 4%. Surprisingly, perhaps, there was no effect of risk

Table 2. Effects of increasing financial risk aversion on management criteria in variance-penalized optimum model.

Criteria	Risk aversion parameter λ: 0	0.0001	0.0002	0.0003	0.0004	0.0005	1.0000
Discounted financial							
VPO ($\$\cdot ha^{-1}$)	1.00	0.70	0.44	0.28	0.17	0.13	0.00
E(NPV) ($\$\cdot ha^{-1}$)	1.00	0.99	0.82	0.62	0.58	0.23	0.00
DNV ($(\$\cdot ha^{-1})^2$)	1.00	0.85	0.56	0.34	0.30	0.06	0.00
var(NPV) ($(\$\cdot ha^{-1})^2$)	1.00	0.91	0.54	0.32	0.28	0.06	0.00
Undiscounted financial							
Annual returns ($\$\cdot ha^{-1}\cdot year^{-1}$)	1.00	0.98	0.82	0.62	0.59	0.18	0.00
Sawlog harvest ($m^3\cdot ha^{-1}\cdot year^{-1}$)	1.00	1.01	0.96	0.82	0.79	0.26	0.00
Pole harvest ($m^3\cdot ha^{-1}\cdot year^{-1}$)	1.00	1.05	1.03	1.05	1.06	0.99	0.00
Total harvest ($m^3\cdot ha^{-1}\cdot year^{-1}$)	1.00	1.03	0.98	0.89	0.88	0.49	0.00
Undiscounted ecological							
Species diversity (index)	1.00	0.97	0.96	0.94	0.94	0.87	1.01
Size diversity (index)	1.00	0.99	1.00	1.00	1.01	1.05	1.11
Basal area ($m^2\cdot ha^{-1}$)	1.00	0.93	0.95	0.98	0.99	1.19	1.70
CO_2e stock ($t\cdot ha^{-1}$)	1.00	0.94	0.95	0.95	0.96	1.04	2.06
Landscape diversity (index)	1.00	0.64	0.73	0.73	0.76	0.81	2.02
Cutting cycle (year)	1.00	0.92	0.95	1.00	1.01	1.42	∞

Note: The criteria values are relative to their level under risk neutrality (λ = 0), shown in Table 1. VPO, variance penalized optimum; E(NPV), expected net present value; DNV, discount normalized value; var(NPV), variance of net present value. Currency: US dollars.

aversion on the diversity of tree size, despite the more frequent harvests that shortened the expected cutting cycle by 5%. The basal area and the stock of CO_2e were 5% to 6% lower, respectively, with risk aversion. Although there was little difference in the total volume of harvest, risk aversion lowered the harvest of sawlogs by 6%, while the harvest of poles increased by 7%.

Variance-penalized optimization

Effects of financial risk aversion

Table 2 shows the results of the second approach to modeling decision making under risk aversion with discounted criteria. This was analogous to the classic certainty-equivalent approach with undiscounted criteria, but substituting the DNV for the variance of NPV.

The variance-penalized optimization model (eq. 12), maximizing the weighted sum of the expected NPV and DNV, was solved with weights of the DNV varying from λ = 0 (risk neutrality) to λ = 0.005 (high risk aversion). The results for λ = 1 are also shown as a limit with absolute risk aversion leading to no harvest. The results in Table 2 are relative to their value in the risk-neutral case (Table 1, column 1).

All four discounted criteria, the variance-penalized optimum (VPO), the expected NPV (*E*(NPV)), the DNV, and the variance of NPV (var(NPV)) decreased monotonically as the risk aversion increased. Increasing risk aversion lowered the variability of the NPV of financial returns at a lower rate than the decrease in variability measured by DNV and var(NPV). There was a close relationship between the DNV and var(NPV) for different λs, which again supported the use of the DNV as a proxy for the variance of the NPV. At the limit for λ = 1, all of the discounted criteria were zero in accordance with the absence of harvest.

Among the undiscounted criteria, the expected annual financial returns decreased with rising financial risk aversion at approximately the same rate as the expected NPV. The corresponding expected annual harvest also decreased with λ, but less rapidly than annual financial returns. Risk aversion substantially decreased the sawtimber harvest, while it had little effect on the pole harvest.

The expected tree species diversity index tended to decrease with moderate increases of λ, but at the limit λ = 1, in the absence of harvest, species diversity was practically the same as with the risk-neutral management. The expected tree size diversity was only raised by high levels of risk aversion, λ = 0.0005, and the extreme λ = 1 with no harvest.

The expected basal area, the stock of CO_2e, and the landscape diversity, in particular, exhibited a U-shaped relationship with risk aversion. They decreased with moderate values of λ and then increased at extreme values. The landscape diversity index decreased by 27% as the risk aversion increased from λ = 0 to λ = 0.002, but it then increased to reach a maximum value under natural conditions in the absence of harvest. A similar U-shaped pattern was observed for the expected length of the cutting cycle, which tended to be lower for moderate risk aversion, but then increased steadily to tend to infinity at λ = 1 in the absence of harvest.

Variance-penalized optimum policies

Table 3 shows the policies that maximized the variance-penalized optimum for the case of risk neutrality, λ = 0, and moderate risk aversion, λ = 0.0002, which, as seen earlier, reduced the variance of the NPV by 54%. For each level of risk aversion, Table 3 shows the stand state after harvest conditional on the stand and market state at decision time. For example, for a stand in state 3 with high basal area only in the hardwood small sawtimber, the best decision of a risk-neutral decision maker was to do nothing, while for a risk-averse decision maker, it was to thin the stand to state 1 by reducing the basal area of hardwood small sawtimber.

Within the 192 possible stand–market states, the decision differed between risk-neutral and risk-averse decision makers for 136 stand–market states: 38 at low market, 46 at medium market, and 52 at high market. Under risk neutrality, 169 stand–market states called for a harvest, 54 at low market, 56 at medium market, and 59 at high market. With risk aversion (λ = 0.0002), the number of stand–market states calling for a harvest increased to 183 (61 at low market, 62 at medium market, and 60 at high market).

As indicated by the coefficients of variation in the last row of Table 1, there was more variation in the postharvest decisions under risk neutrality (λ = 0) than under risk aversion (λ = 0.0002), regardless of the market state.

Summary and conclusion

Since its introduction in Markowitz (1952), the mean-variance analysis has become a fundamental base in financial investment theory. Over six decades later, it still has considerable relevance in research and practice (e.g., Markowitz et al. 2000; De Giorgi and Hens 2009; Fabozzi et al. 2012; García et al. 2015), despite criticisms of its

Table 3. Decisions by stand and market state (low, medium, high) that maximized the variance-penalized discounted value of financial returns for risk-neutral and risk-averse decision makers.

Stand state no.	Stand composition[a]	Decision[b]					
		Risk neutral (λ = 0.0000)			Risk averse (λ = 0.0002)		
		Low	Medium	High	Low	Medium	High
1	000, 000	—	—	—	—	—	—
2	000, 001	1	1	1	1	1	1
3	000, 010	—	—	—	—	—	1
4	000, 011	3	3	3	3	3	3
5	000, 100	1	1	—	1	1	1
6	000, 101	1	1	5	5	5	5
7	000, 110	3	3	3	3	3	5
8	000, 111	3	3	3	7	7	7
9	001, 000	1	1	1	1	1	—
10	001, 001	1	1	1	2	2	9
11	001, 010	3	3	3	3	3	9
12	001, 011	3	3	3	4	4	11
13	001, 100	1	1	5	5	5	9
14	001, 101	1	1	5	6	6	13
15	001, 110	3	3	3	7	7	11
16	001, 111	3	3	3	8	8	15
17	010, 000	—	1	1	1	1	1
18	010, 001	—	—	1	2	2	2
19	010, 010	—	—	3	3	3	3
20	010, 011	19	19	3	19	4	4
21	010, 100	—	—	5	5	5	5
22	010, 101	21	21	5	21	6	6
23	010, 110	19	19	3	7	7	7
24	010, 111	19	19	3	19	8	8
25	011, 000	17	1	1	17	17	9
26	011, 001	18	18	1	18	18	10
27	011, 010	19	19	3	19	19	11
28	011, 011	19	19	3	20	20	12
29	011, 100	21	21	5	21	21	13
30	011, 101	21	21	5	22	22	14
31	011, 110	19	19	3	23	23	15
32	011, 111	19	19	3	24	24	16
33	100, 000	1	1	—	1	1	—
34	100, 001	1	1	33	33	33	33
35	100, 010	3	3	3	3	3	33
36	100, 011	—	—	—	—	35	35
37	100, 100	1	1	33	5	33	33
38	100, 101	1	1	33	2	37	37
39	100, 110	3	3	3	7	35	33
40	100, 111	36	36	36	36	36	35
41	101, 000	—	1	33	33	33	—
42	101, 001	41	1	33	34	34	41
43	101, 010	3	3	3	35	35	41
44	101, 011	36	36	36	36	36	43
45	101, 100	41	1	33	37	37	41
46	101, 101	41	1	33	38	38	41
47	101, 110	3	3	3	39	39	41
48	101, 111	36	36	36	40	40	47
49	110, 000	—	—	33	33	33	33
50	110, 001	18	49	33	49	34	34
51	110, 010	19	19	3	19	35	35
52	110, 011	—	—	36	36	36	36
53	110, 100	21	21	33	21	37	37
54	110, 101	21	21	33	18	38	38
55	110, 110	19	19	3	19	39	39
56	110, 111	52	52	36	52	36	40
57	111, 000	49	49	33	49	49	41
58	111, 001	18	49	33	50	50	42
59	111, 010	19	19	3	51	51	43
60	111, 011	52	52	36	52	52	44

Table 3 (*concluded*).

Stand state no.	Stand composition[a]	Decision[b] Risk neutral (λ = 0.0000) Low	Medium	High	Risk averse (λ = 0.0002) Low	Medium	High
61	111, 100	21	21	33	53	53	45
62	111, 101	21	21	33	54	54	46
63	111, 110	19	19	3	55	55	47
64	111, 111	52	52	36	56	56	48
No. of states with harvest		54	56	59	61	62	60
Total		169			183		
No. of decisions affected by risk aversion					38	46	52
Total					136		
Coefficient of variation		0.88	1.03	1.04	0.84	0.78	0.71

[a]Basal area in pulpwood, small sawtimber, and large sawtimber of pines (first three digits), and hardwoods (last three digits): 1, higher than current average; 0, lower than current average.

[b]Stand state no. resulting from the harvest decision: "—" indicates no harvest.

theoretical limitations and drawbacks in empirical applications (e.g., Steinbach 2001; Michaud 2004).

Within the MDP framework and for undiscounted expected value criteria, it is possible to directly apply mean-variance theory and the closely related certainty-equivalent optimization with quadratic programming. Buongiorno et al. (2017) present examples of these methods for undiscounted expected values of financial and ecological criteria in the context of uneven-aged mixed-species forests in the southern US. For the case of discounted criteria such as the NPV of returns crucial in investment analysis, however, straight mean-variance theory in the MDP framework leads to highly nonlinear optimization problems that are hard to solve in practice and for which convergence is not even guaranteed (Mannor and Tsitsiklis 2011).

In view of this difficulty, as an alternative to the classic variance measure of risk, we adopted the DNV (Filar et al. 1989), the sum of the discounted expected squared deviation of returns with respect to the annualized expected NPV. A strong motivation of the DNV is that the annualized expected NPV can be viewed as a "reference point" for gain or loss as in prospect theory (Tversky and Kahneman 1992). Minimizing the expected fluctuations around the reference point, therefore, is similar to income smoothing (Morduch 1995) and is a reasonable way of modeling objectives such as nonindustrial private landowners' risk-reduction behavior. Furthermore, with the data in this study, there was an almost one to one correspondence between relative changes in the DNV and relative changes in the variance of NPV.

The forests used as the case study are of high commercial importance in the southern US, known as the world's "wood basket" (Oswalt et al. 2014). In this context, management that reduced risk also lowered the NPV of financial returns, in accordance with previous findings (e.g., Eyvindson and Kangas 2016). Management consistent with risk aversion that reduced the variance of the financial NPV by half led to a near 60% lower expected NPV of financial returns. Previous studies have shown that the current actual management generates approximately 25% of the maximum possible NPV (Buongiorno and Zhou 2015). Zhou (2017) attributes this difference between actual and potential revenues to the amenity value of the standing trees. The present results suggest that it may also reflect, in part, the financial risk aversion of the forest owners. How to combine the effects of risk aversion and non-market values in a single model may be a worthwhile topic for future research.

Another notable result of this study was the non-monotonic relationship between financial risk aversion and the expected undiscounted value of nonfinancial criteria. As a case in point, the length of the cutting cycle is an important ecological yardstick for the frequency of harvest-induced disturbances. Although it was found that the expected cutting cycle did decrease with low to moderate levels of financial risk aversion, it then increased steadily with higher financial risk aversion to ultimately approach infinity in the absence harvest. A similar U-shaped effect of financial risk aversion was observed for the expected landscape diversity, the stock of CO_2e stored in forest stands, and the diversity of tree species and size. The higher landscape diversity, at each market level, was induced by observed higher variation of the postdecision stand states under risk neutrality than under risk aversion. Overall, in contrast to previous results with undiscounted financial returns (Buongiorno et al. 2017), low to medium risk aversion of discounted financial returns did not necessarily induce expected undiscounted ecological benefits. The effect of discounting ecological criteria would be worth exploring in future research, with methods similar to those applied here to discounted financial returns.

Acknowledgements

The research leading to this paper was supported in part by a Purdue University Startup fund, National Institute of Food and Agriculture Access No. 1016676, and a joint venture agreement between the University of Wisconsin and the U.S. Department of Agriculture Forest Service Southern Forest Research Station (17-JV-11330143-087), in cooperation with Project Leader Jeff Prestemon.

References

Andersson, M., and Gong, P. 2010. Risk preferences, risk perceptions and timber harvest decisions — an empirical study of nonindustrial private forest owners in Northern Sweden. For. Policy Econ. **12**: 330–339. doi:10.1016/j.forpol.2010.02.002.

Arrow, K.J., Cropper, M.L., Eads, G.C., Hahn, R.W., Lave, L.B., Noll, R.G., Portney, P.R., Russell, M., Schmalensee, R., Smith, V.K., and Stavins, R.N. 1996. Is there a role for benefit–cost analysis in environmental, health, and safety regulation? Science, **272**(5259): 221–222. doi:10.1126/science.272.5259.221. PMID:8602504.

Arrow, K.J., Cropper, M.L., Gollier, C., Groom, B., Heal, G.M., Newell, R.G., Nordhaus, W.D., Pindyck, R.S., Pizer, W.A., Portney, P.R., Sterner, T., Tol, R.S.J., and Weitzman, M.L. 2014. Should governments use a declining discount rate in project analysis? Rev. Environ. Econ. Policy, **8**(2): 145–163. doi:10.1093/reep/reu008.

Buongiorno, J. 2001. Generalization of Faustmann's formula for stochastic forest growth and prices with Markov decision process models. For. Sci. **47**: 466–474.

Buongiorno, J., and Zhou, M. 2015. Adaptive economic and ecological forest management under risk. For. Ecosyst. **2**: 4. doi:10.1186/s40663-015-0030-y.

Buongiorno, J., Zhou, M., and Johnston, C. 2017. Risk aversion and risk seeking in multicriteria forest management: a Markov decision process approach. Can. J. For. Res. **47**(6): 800–807. doi:10.1139/cjfr-2016-0502.

Chung, K.-J., and Sobel, M.J. 1987. Discounted MDP's: distribution functions and exponential utility maximization. SIAM J. Control Optim. **25**: 49–62. doi:10.1137/0325004.

Couture, S., and Reynaud, A. 2011. Forest management under fire risk when forest carbon sequestration has value. Ecol. Econ. **70**: 2002–2011. doi:10.1016/j.ecolecon.2011.05.016.
Couture, S., Cros, M.-J., and Sabbadin, R. 2016. Risk aversion and optimal management of an uneven-aged forest under risk of windthrow: a Markov decision process approach. J. For. Econ. **25**: 94–114. doi:10.1016/j.jfe.2016.08.002.
De Giorgi, E., and Hens, T. 2009. Prospect theory and mean-variance analysis: does it make a difference in wealth management. Investment Manage. Financ. Innovations, **6**(1): 127.
Denardo, E.V. 1967. Contraction mappings in the theory underlying dynamic programming. SIAM Rev. **9**(2): 165–177. doi:10.1137/1009030.
d'Epenoux, F. 1963. A probabilistic production and inventory problem. Manage. Sci. **10**(1): 98–108. doi:10.1287/mnsc.10.1.98.
Eyvindson, K., and Kangas, A. 2016. Integrating risk preferences in forest harvest scheduling. Ann. For. Sci. **73**(2): 321–330. doi:10.1007/s13595-015-0517-2.
Fabozzi, F.J., Markowitz, H.M., Kolm, P.N., and Gupta, F. 2012. Mean-variance model for portfolio selection. *In* Encyclopedia of financial models. *Edited by* F.J. Fabozzi. Wiley Online Library. doi:10.1002/9781118182635.efm0003.
Feliciano, D., Bouriaud, L., Brahic, E., Deuffic, P., Dobsinska, Z., Jarsky, V., Lawrence, A., Nybakk, E., Quiroga, S., Suarez, C., and Ficko, A. 2017. Understanding private forest owners' conceptualisation of forest management: evidence from a survey in seven European countries. J. Rural Stud. **54**: 162–176. doi:10.1016/j.jrurstud.2017.06.016.
Filar, J.A., Kallenberg, L.C., and Lee, H.-M. 1989. Variance-penalized Markov decision processes. Math. Oper. Res. **14**: 147–161. doi:10.1287/moor.14.1.147.
Fleming, W.H., and Hernández-Hernández, D. 1997. Risk-sensitive control of finite state machines on an infinite horizon I. SIAM J. Control Optim. **35**: 1790–1810. doi:10.1137/S0363012995291622.
Fleming, W.H., and Hernández-Hernández, D. 1999. Risk-sensitive control of finite state machines on an infinite horizon II. SIAM J. Control Optim. **37**: 1048–1069. doi:10.1137/S0363012997321498.
Freund, R.J. 1956. The introduction of risk into a programming model. Econometrica, **24**: 253–263. doi:10.2307/1911630.
García, F., González-Bueno, J.A., and Oliver, J. 2015. Mean-variance investment strategy applied in emerging financial markets: evidence from the Colombian stock market. Intellect. Econ. **9**(1): 22–29. doi:10.1016/j.intele.2015.09.003.
Goulder, L.H., and Stavins, R.N. 2002. Discounting: an eye on the future. Nature, **419**(6908): 673–674. doi:10.1038/419673a. PMID:12384678.
Guo, X., Ye, L., and Yin, G. 2012. A mean-variance optimization problem for discounted Markov decision processes. Eur. J. Oper. Res. **220**(2): 423–429. doi:10.1016/j.ejor.2012.01.051.
Hardisty, D.J., and Weber, E.U. 2009. Discounting future green: money versus the environment. J. Exp. Psychol., Gen. **138**(3): 329–340. doi:10.1037/a0016433. PMID:19653793.
Harrison, M. 2010. Valuing the future: the social discount rate in cost-benefit analysis. Visiting Researcher Paper. Australian Government Productivity Commission, Canberra, Australia.
Hillier, F.S., and Lieberman, G.J. 2005. Introduction to operations research. 8th ed. McGraw-Hill, Boston.
Howard, R.A., and Matheson, J.E. 1972. Risk-sensitive Markov decision processes. Manage. Sci. **18**: 356–369. doi:10.1287/mnsc.18.7.356.
Jaquette, S.C. 1976. A utility criterion for Markov decision processes. Manage. Sci. **23**: 43–49. doi:10.1287/mnsc.23.1.43.
Kaya, I., and Buongiorno, J. 1987. Economic harvesting of uneven-aged northern hardwood stands under risk: a Markovian decision model. For. Sci. **33**: 889–907.
Laakkonen, A., Zimmerer, R., Kähkönen, T., Hujala, T., Takala, T., and Tikkanen, J. 2018. Forest owners' attitudes toward pro-climate and climate-responsive forest management. For. Policy Econ. **87**: 1–10. doi:10.1016/j.forpol.2017.11.001.
Lembersky, M.R., and Johnson, K.N. 1975. Optimal policies for managed stands: an infinite horizon Markov decision process approach. For. Sci. **21**: 109–122.
Levitt, S., and Ben-Israel, A. 2001. On modeling risk in Markov decision processes. *In* Optimization and related topics. *Edited by* A. Rubinov and B. Glover. Springer, Boston, Mass. pp. 27–40. doi:10.1007/978-1-4757-6099-6_3.
Lindgren, B.W. 1968. Statistical theory. 2nd ed. MacMillan, London.
Lunnan, A., Nybakk, E., and Vennesland, B. 2006. Entrepreneurial attitudes and probability for start-ups — an investigation of Norwegian non-industrial private forest owners. For. Policy Econ. **8**(7): 683–690. doi:10.1016/j.forpol.2005.06.016.
Mannor, S., and Tsitsiklis, J. 2011. Mean-variance optimization in Markov decision processes. arXiv:1104.5601.
Markowitz, H. 1952. Portfolio selection. J. Finance, **7**: 77–91. doi:10.1111/j.1540-6261.1952.tb01525.x.
Markowitz, H.M., Todd, G.P., and Sharpe, W.F. 2000. Mean-variance analysis in portfolio choice and capital markets. John Wiley & Sons.
Martell, D.L. 1980. The optimal rotation of a flammable forest stand. Can. J. For. Res. **10**(1): 30–34. doi:10.1139/x80-006.
Michaud, R.O. 2004. Why mean-variance optimization isn't useful for investment management. New Frontier Advisors.
Morduch, J. 1995. Income smoothing and consumption smoothing. J. Econ. Perspect. **9**(3): 103–114. doi:10.1257/jep.9.3.103.
Oswalt, S.N., Thompson, M., and Smith, W.B. 2014. US forest resource facts and historical trends. USDA Forest Service, Washington, D.C., For. Serv. FS-1035.
Portney, P.R., and Weyant, J.P. 1999. Discounting and intergenerational equity. Routledge.
Puterman, M.L. 2009. Markov decision processes: discrete stochastic dynamic programming. John Wiley & Sons.
Roessiger, J., Griess, V.C., and Knoke, T. 2011. May risk aversion lead to near-natural forestry? A simulation study. Forestry, **84**: 527–537. doi:10.1093/forestry/cpr017.
Sobel, M.J. 1975. Ordinal dynamic programming. Manage. Sci. **21**(9): 967–975. doi:10.1287/mnsc.21.9.967.
Sobel, M.J. 1982. The variance of discounted Markov decision processes. J. Appl. Probab. **19**: 794–802. doi:10.2307/3213832.
Sobel, M.J. 1994. Mean-variance tradeoffs in an undiscounted MDP. Oper. Res. **42**: 175–183. doi:10.1287/opre.42.1.175.
Steinbach, M.C. 2001. Markowitz revisited: mean-variance models in financial portfolio analysis. SIAM Rev. **43**(1): 31–85. doi:10.1137/S0036144500376650.
Tahvonen, O., and Kallio, M. 2006. Optimal harvesting of forest age classes under price uncertainty and risk aversion. Nat. Resour. Model. **19**: 557–585. doi:10.1111/j.1939-7445.2006.tb00194.x.
Tian, N., Poudyal, N.C., Hodges, D.G., Young, T.M., and Hoyt, K.P. 2015. Understanding the factors influencing nonindustrial private forest landowner interest in supplying ecosystem services in Cumberland Plateau, Tennessee. Forests, **6**(11): 3985–4000. doi:10.3390/f6113985.
Tversky, A., and Kahneman, D. 1992. Advances in prospect theory: cumulative representation of uncertainty. J. Risk Uncertainty, **5**(4): 297–323. doi:10.1007/BF00122574.
Xia, L. 2018. Mean-variance optimization of discrete time discounted Markov decision processes. Automatica, **88**: 76–82. doi:10.1016/j.automatica.2017.11.012.
Zhou, M. 2017. Valuing environmental amenities through inverse optimization: theory and case study. J. Environ. Econ. Manage. **83**: 217–230. doi:10.1016/j.jeem.2017.03.006.
Zhou, M., and Buongiorno, J. 2006. Forest landscape management in a stochastic environment, with an application to mixed loblolly pine–hardwood forests. For. Ecol. Manage. **223**(1–3): 170–182.
Zhou, M., and Buongiorno, J. 2011. Effects of stochastic interest rates in decision making under risk: a Markov decision process model for forest management. For. Policy Econ. **13**: 402–410. doi:10.1016/j.forpol.2011.03.007.