Steven A. Lippman and John J. McCall
University of California, Los Angeles

Wayne L. Winston
Indiana University

Constant Absolute Risk Aversion, Bankruptcy, and Wealth-Dependent Decisions*

We construct a model in which it is rational for an agent with constant absolute risk aversion to select the more risky of two investments if and only if his wealth is small. This paradoxical behavior is induced by the presence of discounting and a bankruptcy constraint, and it bodes ill for the empirical resolution of the controversial assumption of risk-averse agents.

Introduction

The existence of insurance, futures markets, and stock markets indicates the importance of risk aversion in economics. In spite of this, there are few studies which attempt to estimate the degree of an agent's risk aversion (see Friend and Blume 1975).[1] Any study which seeks to quantify risk aversion, regardless of the setting, must necessarily resort to indirect measures. In this article

*This research was partially supported by the National Science Foundation through grant SOC-7808985.

1. This gap is especially significant given the continued controversy regarding the nature of the objective function when agents must produce in an uncertain environment. Some (see Sandmo 1971) have assumed that firms maximize expected utility where utility is a strictly concave function of profits. The model based on the strict concavity assumption gives rise to implications that are quite different from those flowing from expected profit maximization. For example, fixed costs affect output decisions. One objection to this model asserts that firms with strictly concave (or strictly convex) utility functions will not survive when confronted with risk-neutral competitors. In the long run, the average profits of the risk-averse firms will be smaller than those of the risk-neutral firm. Moreover, there will be a tendency for them to disappear or be acquired by risk neutral entrepreneurs (see, e.g., Gould 1976). Furthermore, the demise of the risk-averse entrepreneur will be beneficial for stockholders who will want each firm in their diversified portfolio to maximize expected profits. On the other hand, if there are perfect markets, then Fisher separation obtains and the firm's goal is profit maximization which, empirically, is indistinguishable from a linear utility function. (The empirical work on efficient markets starts with the assumption of perfect capital markets.) This problem continues to be the subject of theoretical inquiry (see Radner 1974).

(*Journal of Business,* 1980, vol. 53, no. 3, pt. 1)

0021-9398/80/5331-0002$01.50

we point out one of the difficulties that will impinge on the validity of empirical work on this topic: indirect measures can be highly sensitive to factors other than risk aversion itself. Just how sensitive the indirect measure can be is illustrated in a paradoxical example in which bankruptcy plays the role of the other "factors."

We consider an individual or firm who seeks to maximize the expected discounted utility *associated* with his income stream where at each point in time he must choose between two risky investments X_1 and X_2. The agent's one-period utility function u is strictly concave (i.e., he is risk averse) and $Eu(w + X_2) > Eu(w + X_1)$ for all wealth levels w so that X_2 is myopically preferred to X_1. Our goal is to construct a simple example wherein the risk-averse agent's optimal strategy is not the myopic one; rather the agent's (dynamic) ordering over the set $\{X_1, X_2\}$ of investment projects is wealth dependent. In particular, the agent chooses investment 1 whenever his wealth falls below ξ and investment 2 whenever his wealth is ξ or more, $1 < \xi < \infty$. In addition, an example is provided in which $E(X_1) > E(X_2)$ so that investment 1 can be thought of as the high risk–high mean return investment. In this case the agent's optimal strategy is truly paradoxical in that he appears to be risk preferent.

In constructing our example, we establish the exact form of the optimal policy (see theorems 2 and 4) as well as quantitative properties of the optimal policy (see theorems 1 and 3); moreover, a simple formula for the optimal return is developed (lemma 1).

The existence of a side constraint (i.e., the possibility of bankruptcy) and the presence of discounting are the sole causes of the anomalous behavior. The point is simply this. Valid inferences concerning an agent's neutrality or aversion to risk must necessarily emanate from a highly robust model. Failure to include a constraint such as bankruptcy might very well produce the maximally incorrect inference.

The Problem

Let X_1 and X_2 be two random variables with cumulative distribution functions F_1 and F_2, u a utility function of the form $u(t) = 1 - e^{-\lambda t}$ with $\lambda > 0$ so that u exhibits constant absolute risk aversion (see Pratt 1964), β the discount factor, $0 \leqslant \beta \leqslant 1$, and $w_0 \geqslant 1$ the agent's initial wealth. Each period the agent observes the size w of his current wealth and selects an investment with distribution F_i; his wealth then becomes $w + X_i$ and he receives the payment $u(w + X_i)$ according to the distribution F_i. (Given i, the draw from F_i is independent of all past draws.) This process repeats indefinitely unless his wealth reaches zero, in which case he is bankrupt (or ruined). Upon becoming bankrupt, the agent's wealth remains zero, whereupon he receives payments of $u(0) = 0$ in each period.

The agent seeks to maximize the expected discounted sum of the utilities of the payments. If the agent's bequest is his wealth at the time of his death, then this is equivalent to maximizing the expected utility of his bequest—with one proviso: the agent's lifetime is a geometric random variable with parameter $1 - \beta$.

In order to secure a model with substantial content, we shall assume: (1) w_0 is an integer and X_1 and X_2 are integer valued, and (2) $F_1(-2) = F_2(-2) = 0$. Assumption 1 assures us that the agent's wealth level is integer valued at all times, whereas assumption 2 precludes the possibility that his wealth might become negative, even after bankruptcy occurs. Together these two assumptions rule out occasions in which the agent's wealth level jumps over the boundary (i.e., zero) rather than landing directly upon it at the time of bankruptcy. If this possibility were not removed, then the one-period utility function u would more properly be defined by $u^*(w) = \max\,[u(w), u(0)]$, all real w, in which case u^* is not concave,[2] though of course its restriction to the interval $[0,\infty)$ is a concave function.

Hakansson (1970) has constructed a finite horizon model with consumption in intermediate periods and constant stochastic returns to scale (as well as several other complicating factors such as labor income) in which "any debt incurred by the individual must at all times be fully secured" (p. 477). There he focuses attention upon the possibility that the multiperiod return function—analogous to $V(\cdot)$ given in our equation (1)—will not be concave even though the (separable) utility of consumption is concave. Our model is designed so that an optimal policy can have positive probability of bankruptcy and our concern is with the optimal strategy rather than the optimal return function.

The Analysis

Standard results in dynamic programming (see Blackwell 1965 and Denardo 1967) permit us to assert that the value $V(w)$ of having initial wealth equal to w is the unique solution[3] to $[V(0) = 0]$:

$$V(w) = \max_{i=1,2} \; [1 - e^{-\lambda w}\mu_i + \beta E V(w + X_i)], \qquad w = 1, 2, \ldots, \quad (1)$$

where $\mu_i \equiv E\bar{e}^{\lambda X_i}$. Moreover, the optimal decision, label it $\pi^*(w)$, is the one that attains the maximum in equation (1). Consequently, the policy $\pi^*(w)$ whenever and however it reaches wealth w, $w = 1, 2, \ldots$, is optimal. Moreover, the standard policy improvement algorithm can be applied to verify the optimality of a given policy (see Orkenyi 1976).[4]

2. This point was made by Professor John P. Gould.

3. If $\beta = 1$, this remains true provided that the expected time till ruin is finite. When $\beta < 1$ there is a uniform bound on the return V_π associated with policy π, as $0 < V_\pi < 1/(1 - \beta)$, so $V(\cdot)$ is bounded.

4. In addition, policy improvement can be utilized when $\beta = 1$ if the expected time till ruin is finite; see the proof of theorem 3.

Henceforth we assume

$$\mu_1 > \mu_2, \tag{2}$$

so that investment 2 is, myopically, preferable to investment 1.

If exclusive employment of investment 2 were to preclude the possibility of bankruptcy, then its myopic domination of investment 1 would suffice to ensure its use at all levels of wealth. This fact is demonstrated in our first theorem (proofs of the theorems and lemmas are to be found in the Appendix), thereby establishing the crucial role of bankruptcy in our model. For convenience, we refer to the policy that always selects investment i as policy I or policy II when $i = 1$ and $i = 2$, respectively.

THEOREM 1: If $F_2(-1) = 0$, then $\pi^* = \text{II}$.

The first step in constructing the example consists of showing that there is a critical number ξ, $1 \leqslant \xi \leqslant +\infty$, such that

$$\pi^*(w) = \begin{cases} 1, & \text{if} \quad w < \xi \\ 2, & \text{if} \quad w \geqslant \xi . \end{cases} \tag{3}$$

We refer to any policy of the form given in equation (3) as a critical number policy.

THEOREM 2: The optimal policy π^* is a critical number policy.[5]

The second and final step is the demonstration that both investments are employed, that is, $\xi > 1$ (1 is used) and $\xi < \infty$ (2 is used). In so doing we need to develop explicit formulas for $V_i(w)$, the return associated with policy i, $i = 1$, II. We begin by defining, for $i = 1, 2$, τ_i to be the first time the agent loses (a net of) \$1.00 when he always selects investment i, $\phi_i(\cdot)$ the generating function of τ_i, and $U_i(w)$ the expected utility associated with policy i when w is the initial wealth and the one-period utility function is $\bar{u}(t) = u(t) - 1 = -e^{-\lambda t}$ for $t > 0$ up to and including the time of bankruptcy and $\bar{u}(0) = 0$ after bankruptcy. (For convenience we shall write $\bar{V}_i = U_i(1)$, $i = 1, 2$.) As shown in Karlin and Taylor (1975, pp. 96–100), ϕ_i is the smallest nonnegative solution of

$$\sum_{k=-1}^{\infty} p_{k,i}\, \phi_i(\beta)^k = 1/\beta, \tag{4}$$

with

$$p_{k,i} = P(X_i = k). \tag{5}$$

In order to keep the analysis simple and to elucidate the intrinsic nature of the process, we shall have occasion to assume that

$$P(X_i = m_i) = p_i > 0 \quad \text{and} \quad P(X_i = -1) = q_i > 0, \qquad i = 1, 2, \tag{6}$$

where $m_1 > m_2 \geqslant 0$. (Because $\mu_1 > \mu_2$ this also implies $q_1 > q_2$.)

5. The proof of theorem 2 is particularly noteworthy as the techniques employed therein have wide applicability.

LEMMA 1: If X_i satisfies equation (6), then the return functions V_i have the following representation (we suppress the subscript i on the right-hand side of eqq. [8] and [9]):

$$V_i(w) = [1 - \phi_i(\beta)]/(1 - \beta) + e^{-\lambda(w-1)}\,\tilde{V}_i + \phi_i(\beta)V_i(w - 1), \tag{7}$$

so that

$$V_i(w) = \frac{1 - \phi(\beta)^w}{1 - \beta} + \tilde{V}e^{-\lambda(w-1)}\frac{1 - [\phi(\beta)e^{\lambda}]^w}{1 - \phi(\beta)e^{\lambda}}, \tag{8}$$

where

$$\tilde{V} = -e^{-\lambda}\mu/\left(1 - p\beta e^{-\lambda m}\left\{\frac{1 - [\phi(\beta)e^{\lambda}]^{m+1}}{1 - \phi(\beta)e^{\lambda}}\right\}\right). \tag{9}$$

Setting $\lambda = 1, \beta = 0.9, m_1 = 2, m_2 = 1, p_1 = 0.5, p_2 = 0.6$, we have $\mu_1 = 1.4268 > 1.3080 = \mu_2$, $\phi_1(0.9) = 0.50953$, $\phi_2(0.9) = 0.48926$, $\tilde{V}_1 = -0.711314$, $\tilde{V}_2 = -0.895855$, so that $V_\pi(1) = 4.22600 > V_2(1) = 4.21148 > 4.19342 = V_1(1)$, where π is the stationary policy that selects investment 2 for all wealth levels except $w = 1$; $V_\pi(1)$ is computed, as per equation (7), as follows:

$$\begin{aligned}V_\pi(1) &= 1 - e^{-\lambda}\mu_1 + p_1\beta V_\pi(3)\\ &= 1 - e^{-\lambda}\mu_1 + p_1\beta\left\{\frac{1 - \phi_2(\beta)}{1 - \beta} + e^{-2\lambda}\tilde{V}_2\right.\\ &\qquad \left. + \phi_2(\beta)\left[\frac{1 - \phi_2(\beta)}{1 - \beta} + e^{-\lambda}\tilde{V}_2 + \phi_2(\beta)V_\pi(1)\right]\right\}\\ &= \frac{1 - e^{-\lambda}\mu_1 + p_1\beta\left\{\dfrac{1 - \phi_2(\beta)^2}{1 - \beta} + \tilde{V}_2[e^{-2\lambda} + e^{-\lambda}\phi_2(\beta)]\right\}}{1 - p_1\beta\phi_2(\beta)^2}.\end{aligned}$$

Because $V_\pi(1) > V_2(1) > V_1(1)$, we know that π^* is neither I nor II; equivalently, π^* employs both investment 1 and 2 so $1 < \xi < \infty$, as desired.[6] Moreover, investment 1 is the high risk–high mean return investment, for $\mu_1 > \mu_2$ and $E(X_1) = 0.5 > 0.2 = E(X_2)$, thereby establishing the existence of the paradoxical behavior as promised.

Finally, the following lemma reveals that the return function $V(\cdot)$ is concave in this example;[7] thus, the paradoxical behavior is not induced by a lack of concavity in the return function.

LEMMA 2: If $F_2(t + 1) \leq F_1(t)$ for all t, then the return function $V(\cdot)$ is concave.

6. It is interesting to note that investment 1 is utilized even though $[1 - \phi_1(0.9)]/[1 - \beta] < [1 - \phi_2(0.9)]/[1 - \beta]$, i.e., investment 1 has the shorter discounted length of play. Evidently, the saving grace of investment 1 rests upon the fact that it has the smaller ruin probability, i.e., $\phi_1(1) = 0.579 < 0.592 = \phi_2(1)$.

7. The concavity of $V_i(\cdot)$ is easily established by inspection of eq. (8), $\phi_i(\beta)e^{\lambda} < 1$, and $\tilde{V}_i < 0$.

The Importance of $\beta < 1$

It was stated in the Introduction that the sole causes of the anomalous behavior are the presence of discounting and the possibility of bankruptcy. The essential role of discounting is demonstrated in the following theorem. (Note that we do not assume X_i is given by eq. [6].)

THEOREM 3: If $\beta = 1$ and $E(X_1) > E(X_2)$, then $\pi^* = I$; that is, if there is no discounting, then it is always optimal to select investment 1.[8]

Because $\mu_1 > \mu_2$, the agent would "prefer" investment 2 on each given play if there were no ramifications beyond the fact that it has the larger one-period utility. However, when $\beta = 1$ this is more than compensated by the fact that the expected time till ruin is larger for investment 1.[9] But the fact that investment 1 prolongs the expected duration of play ($= wE\tau_i$) does not imply that it has the greater expected discounted number of plays ($= [1 - \phi_i(\beta)^w]/(1 - \beta)$) till ruin. This "explains" why $\xi < \infty$ in our example when $\beta < 1$ and $\xi = \infty$ when $\beta = 1$.

The Option of Not Investing[10]

In light of theorem 1, it is natural to inquire whether the behavior encapsulated by equation (3) persists if the agent is allowed to stop altogether. To wit, suppose that in addition to the risky investments X_1 and X_2 the riskless investment $X_0 \equiv 0$ is also available to the agent each period. As before, $\mu_1 > \mu_2$, and in addition we assume $\mu_0 > \mu_1$, so both risky investments myopically dominate "not investing." We now provide a numerical example in which the optimal policy satisfies equation (3) with $\xi > 1$ and $\xi < \infty$; that is, the wealth-dependent (dynamic) ordering of the investment projects persists and the agent never selects investment 0. Toward this end, we establish a version of theorem 2 to account for the possibility of using investment 0. It asserts that there is no reason for selecting investment 0.

THEOREM 4: The optimal policy π^* is a critical number policy, that is, it satisfies equation (3); in particular, investment 0 is not utilized.

To obtain the desired numerical example, we set $\lambda = 1$, $\beta = 0.4$, $m_1 = 1$, $m_2 = 2$, $p_1 = 0.746$, and $p_2 = 0.68$. Then $1 > \mu_1 = 0.96488 > 0.96188 = \mu_2$, $\phi_1(0.4) = 0.10488$, $\phi_2(0.4) = 0.12858$, $\tilde{V}_1 = -0.41326$, $\tilde{V}_2 = -0.37412$, so that

$$V_\pi(1) = 1.07875 > V_1(1) = 1.07860 > V_2(1) = 1.07825, \quad (10)$$

8. The proof reveals that always selecting investment 1 is uniquely optimal among the class of stationary policies if $E(X_2) < 0$.

9. If $E(X_2) < 0$, then from Wald's equation and $E(X_1) > E(X_2)$ we have $E\tau_1 = -1/E(X_1) > -1/E(X_2)$; if $E(X_1) > E(X_2) \geqslant 0$, then $E\tau_1 = E\tau_2 = \infty$.

10. This line of inquiry was suggested by the referee.

where π is the stationary policy that selects investment 2 for all wealth levels except $w = 1$. As per equation (7), we have

$$V_\pi(1) = \left\{1 - e^{-\lambda}\mu_1 + p_1\beta\left[\frac{1 - \phi_2(\beta)}{1 - \beta} + e^{-\lambda}\hat{V}_2\right]\right\} / [1 - p_1\beta\phi_2(\beta)].$$

Equation (10) reveals that π^* is neither I nor II, as desired.

Unlike our previous numerical example, investment 1 is the low-return investment. Whereas the (dynamic) wealth-dependent ordering of the investments is preserved, the anomaly disappears.[11]

Appendix

PROOF OF THEOREM 1: Let $< A_i >$ and $< B_i >$ be independent sequences of i.i.d. random variables with distribution F_1 and F_2, respectively, and set $S_n = \Sigma_{i=1}^n B_i$. Because $F_2(-1) = 0$, it is clear that

$$\begin{aligned} V_2(w) &= E\sum_{n=1}^{\infty}\beta^{n-1}[1 - e^{-\lambda(w+S_n)}] = \frac{1}{1-\beta} \\ &\quad - e^{-\lambda w}\sum_{n=1}^{\infty}\beta^{n-1}\, Ee^{-\lambda S_n} \\ &= \frac{1}{1-\beta} - e^{-\lambda w}\sum_{n=1}^{\infty}\beta^{n-1}\mu_2^n = \frac{1}{1-\beta} \\ &\quad - e^{-\lambda w}\frac{\mu_2}{1-\beta\mu_2}. \end{aligned}$$

Let π be the nonstationary policy that selects investment 1 in some state w, say $\hat{w}$, and thereafter always selects investment 2. Then for $\hat{w} > 1$ we have

$$\begin{aligned} V_\pi(\hat{w}) &= E\sum_{n=1}^{\infty}\beta^{n-1}[1 - e^{-\lambda(\hat{w}+A_1+S_{n-1})}] \\ &= \frac{1}{1-\beta} - e^{-\lambda\hat{w}}\frac{\mu_1}{1-\beta\mu_2} < V_2(\hat{w}), \end{aligned}$$

for $\mu_1 > \mu_2$ and $\mu_2 \leqslant 1$ because $F_2(-1) = 0$.

11. Unfortunately the condition $E(X_1) < E(X_2)$ is necessary. The reason for this is as follows. If $\mu_1 = \mu_2$ and $E(X_1) > E(X_2)$, then $\phi_1(1) - \phi_2(1) > (=) [<] 0$ whenever $\mu_2 < (=) [>] 1$. This fact is most easily discerned from the observations, first, that $\pi_1 \equiv \phi_1(1) = e^{-\lambda} = \phi_2(1) \equiv \pi_2$ when $\mu = 1$, where μ is the common value of the μ_i; and, second, $d\pi_2/dp_2 = (1 - \pi_2^2)/(2p_2\pi_2 - 1)$ and $d\pi_1/dp_2 = (1 - \pi_1^3)/(3p_2\pi_1^2 - 1)$, where $\gamma = (e^\lambda - e^{-\lambda})/(e^\lambda - e^{-2\lambda})$. Combining these with the fact that $p_2 = (e^\lambda - 1)/(e^\lambda - e^{-\lambda})$ when $\mu = 1$ reveals that $d\pi_2/dp_2\,|_{\mu=1} < d\pi_1/dp_2\,|_{\mu=1}$, whence $\pi_2 < \pi_1$ when $\mu < 1$. In particular, when $1 > \mu_1 > \mu_2$ and $E(X_1) > E(X_2)$, we necessarily must have $\phi_1(\beta) > \phi_2(\beta)$ for all $\beta \in [0,1]$; hence, investment 1 has the smaller discounted length of play as well as the smaller return per play—in short, it has nothing to recommend its use.

Next, notice that

$$E \sum_{n=1}^{\infty} \beta^{n-1}(1 - e^{-\lambda S_n}) = (1 - \mu_2)/(1 - \beta)(1 - \beta\mu_2)$$

$$\geq 0 = V_2(0).$$

Consequently,

$$\begin{aligned} V_\pi(1) &= 1 - \mu_1 e^{-\lambda} + \beta E V_2(1 + A_1) \\ &\leq 1 - \mu_1 e^{-\lambda} + \beta e^{-\lambda} E \sum_{n=1}^{\infty} \beta^{n-1}[1 - e^{-\lambda(S_n + A_1)}] \\ &= \frac{1}{1-\beta} - e^{-\lambda} \frac{\mu_1}{1 - \beta\mu_2} < V_2(1). \end{aligned}$$

Thus II cannot be improved so the policy-improvement results of Orkenyi (1976) imply that II is optimal. Q.E.D.

PROOF OF THEOREM 2: If π^* is not of the form given in equation (3), then there is an integer $\hat{w}$ with $\hat{w} \geq 2$ such that $\pi^*(\hat{w} - 1) = 2$ and $\pi^*(\hat{w}) = 1$.

Define the random variable $Q(\geq 1)$ to be the number of times that π^* consecutively selects investment 1 when starting from the state $\hat{w}$, and let π be the nonstationary policy which selects investment 2 at time 1 and investment 1 at times $2, 3, \ldots, Q + 1$, and then chooses investment $\pi^*(w)$ when in state w. Let $< A_i >$ and $< B_i >$ be independent sequences of i.i.d. random variables with distribution F_1 and F_2, respectively. We take A_i (B_i) to be the outcome the ith time investment 1 (2) is selected. As defined, Q is a stopping time with respect to $\{A_i\}$.

By construction, the state of the system is, with probability 1, the same after time $Q + 1$ when using either π or π^* and starting at $\hat{w}$; moreover, $1 - \beta^Q \exp [-\lambda(\hat{w} + B_1 + \Sigma_1^Q A_i)]$ is the reward earned at time $Q + 1$ under both policies. Consequently, the fact that $\{Q > i\}$ and $\exp(-\lambda \Sigma_{j=1}^i A_j)$ are measurable with respect to the σ-algebra generated by $\{A_1, \ldots, A_i\}$ enables us to write, without concern for whether or not $P(Q < \infty) = 1$,

$$\begin{aligned} e^{\lambda\hat{w}}[V_\pi(\hat{w}) - V(\hat{w})] &= E \sum_{i=0}^{Q-1} \beta^i \left(e^{-\lambda \sum_{j=1}^{i+1} A_j} - e^{-\lambda \sum_{j=1}^{i} A_j - \lambda B_1} \right) \\ &= \sum_{i=0}^{\infty} \beta^i E \left[1_{(Q>i)} e^{-\lambda \sum_{j=1}^{i} A_j} (e^{-\lambda A_{i+1}} - e^{-\lambda B_1}) \right] \\ &= \sum_{i=0}^{\infty} \beta^i E \left\{ E \left[1_{(Q>i)} e^{-\lambda \sum_{j=1}^{i} A_j} (e^{-\lambda A_{i+1}} - e^{-\lambda B_1}) \,\middle|\, A_1, \ldots, A_i \right] \right\} \\ &= \sum_{i=0}^{\infty} \beta^i E \left[1_{(Q>i)} e^{-\lambda \sum_{j=1}^{i} A_j} E(e^{-\lambda A_{i+1}} - e^{-\lambda B_1} \mid A_1, \ldots, A_i) \right] \end{aligned}$$

$$= (\mu_1 - \mu_2) \sum_{i=0}^{\infty} \beta^i E \left[1_{\{Q>i\}} e^{-\lambda \sum_{j=1}^{i} A_j} \right]$$

$$= (\mu_1 - \mu_2) E \sum_{i=0}^{Q-1} \beta^i e^{-\lambda \sum_{j=1}^{i} A_j} > 0,$$

as $\mu_1 - \mu_2 > 0$ and each term in the sum is positive. This contradicts the optimality of π^*. Q.E.D.

PROOF OF LEMMA 1: Suppressing the subscript i, it is clear that

$$\tilde{V} \equiv U(1) = -e^{-\lambda}\mu + \beta p U(m + 1). \tag{A1}$$

Now $U(k + 1)$ can be expressed as the sum of two components. The first is $e^{-\lambda k}\tilde{V}$, the value of the payments received during the first τ time periods, that is, until the agent's wealth falls by \$1.00 for the first time. The second component is $\phi(\beta)\, U(k)$, where $\phi(\beta) = E\beta^\tau$ is the expected discounted time until the agent loses (a net of) \$1.00. Thus,

$$U(k + 1) = e^{-\lambda k}\tilde{V} + \phi(\beta)U(k). \tag{A2}$$

By repeated application of equation (A2) to (A1), we obtain, as per equation (9),

$$\begin{aligned}
\tilde{V} &= -e^{-\lambda}\mu + \beta p[e^{-\lambda m}\tilde{V} + \phi(\beta)U(m)] \\
&= -e^{-\lambda}\mu + \tilde{V}\beta p e^{-\lambda m} + \beta p \phi(\beta)[e^{-\lambda(m-1)}\tilde{V} + \phi(\beta)U(m - 1)] \\
&= \ldots = -e^{-\lambda}\mu + \tilde{V}\beta p e^{-\lambda m} \sum_{k=0}^{m} [\phi(\beta)e^{\lambda}]^k \\
&= -e^{-\lambda}\mu \Big/ \left(1 - \beta p e^{-\lambda m}\left\{\frac{1 - [\phi(\beta)e^{\lambda}]^{m+1}}{1 - \phi(\beta)e^{\lambda}}\right\}\right).
\end{aligned}$$

Writing $u(t) = 1 + \tilde{u}(t)$, and employing the argument used to develop equation (A2) it is clear that

$$V_i(w) = [1 - \phi(\beta)]/(1 - \beta) + e^{-\lambda(w-1)}\tilde{V} + \phi(\beta)V_i(w - 1), \tag{A3}$$

as $E \sum_{n=1}^{\tau} \beta^{n-1} = E\{(1 - \beta^\tau)/(1 - \beta) = [1 - \phi(\beta)]/(1 - \beta)\}$. Iterating equation (A3) yields

$$\begin{aligned}
V_i(w) &= \frac{1 - \phi(\beta)}{\text{e}1 - \beta} \sum_{k=1}^{w} \phi(\beta)^{k-1} + e^{-\lambda(w-1)}\tilde{V} \sum_{k=1}^{w} [e^{\lambda}\phi(\beta)]^{k-1} \\
&= \frac{1 - \phi(\beta)}{1 - \beta}\, \frac{1 - \phi(\beta)^w}{1 - \phi(\beta)} + e^{-\lambda(w-1)}\tilde{V}\, \frac{1 - [e^{\lambda}\phi(\beta)]^w}{1 - e^{\lambda}\phi(\beta)}.
\end{aligned}$$

Q.E.D.

PROOF OF LEMMA 2: Let V_n be the n-period analogue of V with $V_0 \equiv 0$. The proof of theorem 2 establishes that $\pi_n(w)$, the optimal action when n periods remain and the current wealth is w, satisfies equation (3). As $V_1(w) = 1 - e^{-\lambda w}$ μ_2, V_1 is concave.

Assume V_n is concave, and define μ_w to be μ_1 or μ_2 as $\pi_{n+1}(w) = 1$ or 2. Similarly, let the random variable Y_w have distribution F_1 or F_2 as $\pi_{n+1}(w) = 1$ or 2, $w = 1, 2, \ldots$. Now fix w. Because $\pi_{n+1}(w)$ and $\pi_{n+1}(w + 2)$ are both suboptimal from state $w + 1$, we have

$$\begin{aligned}\Delta_{n+1}(w) &\equiv V_{n+1}(w + 1) - V_{n+1}(w) - [V_{n+1}(w + 2) - V_{n+1}(w + 1)]\\ &\geqslant (1 - e^{-\lambda})e^{-\lambda w}(\mu_w - e^{-\lambda}\mu_{w+2})\\ &\quad + \beta E\{[V_n(w + 1 + Y_w) - V_n(w + Y_w)]\\ &\quad - [V_n(w + 2 + Y_{w+2}) - V_n(w + 1 + Y_{w+2})]\}.\end{aligned}$$

Either $\mu_w = \mu_{w+2}$ and $Y_w = Y_{w+2}$ or $\mu_w = \mu_1 > \mu_2 = \mu_{w+2}$ and Y_w has distribution F_1 and Y_{w+2} has distribution F_2. In the former case the concavity of V_n implies that $\Delta_{n+1}(w) \geqslant 0$, whereas $\mu_1 > \mu_2$, the concavity of V_n, and the hypothesis $Y_{w+2} + 1$ stochastically larger than Y_w yields $\Delta_{n+1}(w) \geqslant 0$ in the latter case. This completes the induction argument.

With $\beta < 1$, $V = \lim_{n\to\infty} V_n$ whereby V inherits the concavity of $< V_n >$. Q.E.D.

Proof of Theorem 3: We assume $E\tau_1 < \infty$, for otherwise $V_1(w) \equiv +\infty$ and I is optimal. The first step in verifying the optimality of I is to develop a formula for the return functions $V_i(\cdot)$; we proceed as in the proof of lemma 1. To begin, define A by

$$A = \sum_{i=0}^{\infty} a_i, \tag{A4}$$

where

$$a_i = \sum_{k=0}^{\infty} P_{k+i}e^{-\lambda k}, \qquad i = 0, 1, 2, \ldots, \tag{A5}$$

and observe that $P(\tau < \infty) = 1$ enables us to assert that the analogue of equation (A2) holds:

$$U(k + 1) = e^{-\lambda k}\hat{V} + U(k), \qquad \text{for } k \geqslant 1. \tag{A6}$$

Furthermore, the analogue of equation (A1) holds and we can iterate (A6) as follows:

$$\begin{aligned}\hat{V} &= -e^{-\lambda}\mu + \sum_{k=0}^{\infty} p_k U(k + 1) = -e^{-\lambda}\mu + \hat{V}p_0 + \sum_{k=1}^{\infty} p_k[e^{-\lambda k}\hat{V} + U(k)]\\ &= -e^{-\lambda}\mu + \hat{V}\sum_{k=0}^{\infty} p_k e^{-\lambda k} + p_1\hat{V} + \sum_{k=1}^{\infty} p_{k+1}U(k + 1)\\ &= -e^{-\lambda}\mu + \hat{V}a_0 + p_1\hat{V} + \hat{V}\sum_{k=1}^{\infty} p_{k+1}e^{-\lambda k} + \sum_{k=1}^{\infty} p_{k+1}U(k)\\ &= -e^{-\lambda}\mu + \hat{V}a_0 + \hat{V}a_1 + p_2\hat{V} + \sum_{k=1}^{\infty} p_{k+2}U(k + 1)\end{aligned}$$

$$= \ldots = -e^{-\lambda}\mu + \hat{V}\sum_{i=0}^{n-1} a_i + \sum_{k=0}^{\infty} p_{k+n}U(k+1)$$

$$= \ldots = -e^{-\lambda}\mu + \hat{V}A,^{12}$$

so that

$$\hat{V} = -e^{-\lambda}\mu/(1-A). \tag{A7}$$

Moreover,

$$(1-e^{-\lambda})A = 1 - \mu e^{-\lambda}. \tag{A8}$$

To see this, simply rearrange the sum of nonnegative terms:

$$A = \sum_{i=0}^{\infty} a_i = \sum_{i=0}^{\infty}\sum_{k=0}^{\infty} p_{k+1}e^{-ik}$$

$$= \sum_{j=0}^{\infty} p_j \sum_{n=0}^{j} e^{-\lambda n} = \sum_{j=0}^{\infty} p_j[1 - e^{-\lambda(j+1)}]/(1-e^{-\lambda})$$

$$= \left[1 - p_{-1} - e^{-\lambda}\left(\sum_{j=0}^{\infty} p_j e^{-\lambda j} + p_{-1}e^{\lambda} - p_{-1}e^{\lambda}\right)\right]/(1-e^{-\lambda}).$$

Finally, it is also clear that

$$V_i(w) = wE\tau + \hat{V}\frac{1-e^{-\lambda w}}{1-e^{-\lambda}}, \tag{A9}$$

for

$$V_i(w) = E\tau + \hat{V}e^{-\lambda(w-1)} + \hat{V}_i(w-1)$$

$$= wE\tau + \hat{V}[e^{-\lambda(w-1)} + e^{-\lambda(w-2)} + \ldots + 1].$$

Let π be the nonstationary policy that selects investment 2 in some state w, say $w = \hat{w}$, and thereafter always selects investment 1. From equation (A9) and the definition of π we can conclude that

$$V_1(\hat{w}) - V_\pi(\hat{w}) = 1 - \mu_1 e^{-\lambda\hat{w}} - (1 - \mu_2 e^{-\lambda\hat{w}})$$

$$+ \sum_{k=-1}^{\infty} V_1(\hat{w}+k)(p_{1,k} - p_{2,k})$$

12. Noting that

$$S_n = \sum_{i=1}^{n} A_i \geq -1 \text{ for } n = 1, 2, \ldots, \tau_1,$$

we have

$$-U(k+1) = E\sum_{n=1}^{\tau_1} e^{-\lambda(k+1+S_n)} - U(k) \leq e^{-\lambda k}E\tau_1 - U(k)$$

$$\leq \ldots \leq E\tau[1 - e^{-\lambda(k+1)}]/(1-e^{-\lambda})$$

so that

$$0 \leq -\sum_{k=0}^{\infty} p_{k+i}U(k+1) \leq E\tau_1 \sum_{k=0}^{\infty} p_{k+i}[1-e^{-\lambda(k+1)}]/(1-e^{-\lambda})$$

$$\leq E\tau_1 \sum_{j=i}^{\infty} p_j/(1-e^{-\lambda}) \to 0,$$

justifying the limiting operation implicit here.

$$= (\mu_2 - \mu_1)e^{-\lambda\hat{w}} + E\tau_1[E(X_1) - E(X_2)]$$
$$+ \frac{\tilde{V}_1}{1 - e^{-\lambda}} \sum_{k=-1}^{\infty} [1 - e^{-\lambda(\hat{w}+k)}](p_{1,k} - p_{2,k})$$
$$= (\mu_1 - \mu_2)e^{-\lambda\hat{w}} \left\{\frac{-\tilde{V}_1}{1 - e^{-\lambda}} - 1\right\} + E\tau_1[E(X_1) - E(X_2)].$$

By definition we must have $\tilde{V} < 0$, so equation (A7) reveals that $A < 1$. Consequently, coupling $\mu_1 > \mu_2$, $E(X_1) > E(X_2)$, $A < 1$, (A7), and (A8) yields $V_1(\hat{w}) - V_\pi(\hat{w}) > 0$, so I is unimprovable. Now $V(w) \leq wE\tau_1 < \infty$ and $EV(w + \Sigma_{i=1}^n B_i) \leq E\tau_1 E[(w + \Sigma_{i=1}^n B_i)^+] \to 0$ as $n \to \infty$; thus policy improvement works in our countable-state Markov decision process (see theorem 2.5 of Orkenyi [1976]) and $\pi^* = I$ is optimal. Q.E.D.

Proof of Theorem 4: Fix $w \geq 1$ and observe that if π is a nonstationary policy that selects investment 2 in state w and thereafter always selects investment 0, and if $\hat{\pi}$ is any stationary policy that selects investment 0 when in state w, then (even if $w = 1$)

$$V_\pi(w) = 1 - e^{-\lambda w}\mu_2 + \frac{\beta}{1 - \beta} \sum_{i=-1}^{\infty} p_{2,i}[1 - e^{-\lambda(w+i)}]$$
$$= 1 - e^{-\lambda w}\mu_2 + \frac{\beta}{1 - \beta}(1 - e^{-\lambda w}\mu_2)$$
$$= (1 - e^{-\lambda w}\mu_2)/(1 - \beta) > (1 - e^{-\lambda w})/(1 - \beta)$$
$$= V_{\hat{\pi}}(w),$$

as $\mu_2 < \mu_0 = 1$. Thus, any policy which employs investment 0 can be improved upon. Consequently, the possibility of using investment 0 can be ignored and the proof of theorem 2 can be utilized. Q.E.D.

References

Blackwell, D. 1965. Discounted dynamic programming. *Annual of Mathematical Statistics* 36:226–35.

Denardo, E. 1967. Contraction mappings in the theory underlying dynamic programming. *SIAM Review* 9:165–77.

Friend, I., and Blume, M. 1975. The demand for risky assets. *American Economic Review* 65:900–922.

Gould, J. P. 1976. Rational expectations and the theory of the firm under uncertainty. Chicago: University of Chicago, Center for Mathematical Studies in Business and Economics.

Hakansson, N. 1970. Friedman-Savage utility functions consistent with risk aversion. *Quarterly Journal of Economics* 84:472–87.

Karlin, S., and Taylor, H. M. 1975. *A First Course in Stochastic Processes*. 2d ed. New York: Academic Press.

Orkenyi, P. 1976. Optimal control of semi-Markovian queueing systems. Stanford, Calif.: Stanford University, Department of Operations Research.

Pratt, J. W. 1964. Risk aversion in the small and in the large. *Econometrica* 32:122–36.

Radner, R. 1974. A note on the unanimity of stockholders' preferences among alternative production plans: a reformulation of the Ekern-Wilson model. *Bell Journal of Economics* 5:181–86.

Sandmo, A. 1971. On the theory of the competitive firm under price uncertainty. *American Economic Review* 61:65–73.