Optimal Phosphorus Loading for a Potentially Eutrophic Lake

Donald Ludwig

Departments of Mathematics and Zoology

University of British Columbia, Vancouver, B. C. V6T 1Z2

Tel. 604-541-9409. Email: ludwig@math.ubc.ca

Steve Carpenter

Center for Limnology, 680 North Park Street

University of Wisconsin, Madison Wisconsin 53706 USA

Email: srcarpen@facstaff.wisc.edu

W. Brock

Department of Economics

University of Wisconsin, Madison Wisconsin 53706 USA

Email: brock@macc.wisc.edu

May 8, 2002

ABSTRACT

We calculate phosphorus (P) loadings for a potentially eutrophic lake that optimize the expected discounted net benefits. The benefits accrue to agricultural interests from activities that result in loading and costs accrue to other interests from the resulting deterioration of water quality. We extend earlier results in Carpenter, Ludwig and Brock (1999) to account for dependence of P recycling upon the concentration of P in sediments.

We find a strong interaction between economic and ecological parameters in determining the policy: the economic discount rate determines whether the time horizon is long or short, and this in turn strongly influences the magnitude of the optimal loadings. Simple policies that neglect dynamics of P in the sediments are inadequate unless the time horizon is short and the dynamics are slow. A stochastic model is essential if there are substantial random fluctuations in loadings. Similarly, uncertainty in the determination of the critical P density that triggers recycling cannot be neglected. Our results may be interpreted as a quantitative precautionary principle that takes account of both economic and ecological aspects of the problem. Our results may also be used to illustrate ideas such as sustainable development, natural capital, and income.

KEY WORDS

phosphorus recycling, optimal phosphorus loading, precautionary principle, slow dynamics, uncertainty and fluctuations, wealth and sustainability

1 Introduction

According to a commonly accepted paradigm, ecosystems are organized according to a nested array of spatial and temporal scales (Stommel 1963, O'Neill et al. 1986, Levin 1992). It follows that one should not attempt to anticipate the future condition of an ecosystem from processes at a single spatial-temporal scale. For example, the dynamics of a hectare of forest over the next 10 years may depend on regional processes of soil formation over the past millennium, the composition of the surrounding mosaic of forest patches for the past century, and population dynamics of herbivorous insects which turn over in weeks at the spatial scale of twigs. It is a major challenge to develop management tools that cope with the multiscalar nature of ecosystems (Carpenter and Turner, 2000).

Interactions of slow and fast variables can create resilient ecosystems or vulnerable ones (Holling 1973). Resilient ecosystems vary within some bounds despite certain types of external shocks to the system. An example is the mosaic structure of forests which is maintained through regular disturbance by wind throw, fire, or insect outbreak and will exist within statistically-definable bounds over a broad range of stochastic regimes (Turner, Gardner and O'Neill, 2001). Vulnerable ecosystems, on the other hand, may change substantially after shocks without returning to anything like the condition that existed prior to the shock. The lake trout (*Salvelinus namaycush*) populations of the Laurentian Great Lakes of North America (Christie 1974) provide an example. Large adult lake trout can survive attack by sea lamprey (*Petromyzon marinus*), but smaller adult lake trout are killed. Fishing, by selectively removing large lake trout, left the populations vulnerable to invasion of the Great Lakes by sea lamprey, which quickly extirpated the lake trout from all of the Great Lakes except Lake Superior. Had fishing targeted smaller individuals and conserved larger ones of this long-lived species, the populations may very well have been resilient to sea lamprey invasion.

Is it possible to manage ecosystems in ways that maintain, or even enhance, resilience? To approach this question using mathematical models, one must employ models that represent the multi-scale nature of ecosystems. Modelers have addressed management in diverse ways, such as simulation of scenarios, creation of "flight simulators" to expose managers to situations they might encounter in real ecosystems, and optimal control analyses. Optimal control is the "gold standard" in economic analysis, but has rarely (if ever) been applied to multi-scale ecological models.

We determine optimal strategies for management of a lake that take account of both the benefits to agriculture of such loadings and the loss of amenity as a consequence of eutrophication. This system exhibits multiple equilibria, a threshold, and several temporal scales. Such multiple equilibria and multiple time scales are a common concern and this work has implications for both ecosystem management and economic theory. We extend the work in Carpenter, Ludwig and Brock (1999) to account for variation of internal loading due to changing levels of dissolved P in the sediments. Our analysis focuses on two state variables: the phosphorus in the water and the phosphorus in the lake sediment. Our objective is to find the optimal loading as a function of our two state variables. The discounted present value of the optimal strategy depends upon the initial values of these variables.

1.1 Lake ecology

Lakes and their watersheds have been an icon for ecosystem analysis since the dawn of the science (Golley 1993). Eutrophication is a state of lakes characterized by toxic blooms of algae, episodes of anoxia, fish kills, and loss of freshwater ecosystem services such as drinking water, water for irrigation and industrial use, fisheries, and recreation (Postel and Carpenter 1997). The fundamental cause is enrichment with phosphorus (P) (Schindler 1977). P is an element in agricultural fertilizers and livestock feed that tends to accumulate in soil (Figure 1). Erosion washes the soil into streams and lakes, where it dissolves and becomes available to algae, including toxic bloom-forming species which are favored at high P concentrations. P also accumulates in sediments. It is rapidly recycled from sediments during episodes of anoxia, which are caused by decay of dead algae.

A minimal model of P dynamics in lakes (Figure 1) includes processing of P by outflow and sedimentation (straight line), and inputs plus recycling (curved lines). Soil P affects the probability distribution of input events to the lake. These input events are driven by variable and unpredictable storms and snow melt, so we model them as a random variable. The logarithm of inputs usually fits a Student-t or gamma distribution (Carpenter et al. 1999a, Reed-Anderson 1999).

Utilities derive from activities that cause P loading to the lake, such as intensive use of fertilizers or animal-intensive farming. These utilities increase (with decreasing slope) as P loading rate increases. Utilities are also associated with water quality as measured by the amount of P in the lake water. These utilities are ecosystem services that may not have markets, such as water for diverse human uses, the state of human health as affected by algal blooms, fisheries, or recreation (Postel and Carpenter 1997). These utilities increase with P at low P levels because a small amount of P supports production of ecosystem services such as fishes and waterfowl. Above an optimal level of P in the lake, however, utility declines.

We do not treat the dynamics of soil phosphorus explicitly. Instead, the effect of soil phosphorus is represented by the mean annual loading rate of phosphorus to the lake. The dynamics are stochastic via shocks to the loading. A one dimensional state variable, P_t, "phosphorus sequestered in algae," generates social disamenities each period which are proportional to its square. Loadings L_t generate private profits proportional to L_t. We assume that utility per period is given by $U_t = \alpha L_t - \beta_1 P_t - \beta_2 {P_t}^2$, for some constants $\alpha > 0$, $\beta_1, \beta_2 \geq 0$. If $P_t > P_c$ where P_c is a "critical" level, recycling of P occurs and P is added to the lake from sedimented phosphorus whose stock we call "mud" M_t. The amount of this recycling in period t is proportional to M_t. Hence the state variable P_t is driven by another state variable M_t in our model. The managing agency is assumed to act in the overall social interest and regulates loadings L_t to maximize the expectations of the discounted sum of utilities.

1.2 Economic interpretation

The economics literature on infinite horizon optimal control models with nonconvexities in the dynamics is large. Here we only review the most recent pieces which have computational content. Beyn et al. (2000), Deissenberg et al. (2000), Dechert and Brock (2000), Moran and Maroto (2002), and Semmler and Sieveking (1999) have studied closely related computational problems. All except Moran and Maroto study deterministic problems with nonconcave dynamics. Moran and Maroto study stochastic one dimensional state variable problems with nonconcave utility and nonconcave dynamics. They provide a detailed theoretical analysis of computational methods for their problem as well as some computed examples. There is a large and growing literature in computational economics on stochastic multidimensional state variable problems but the vast bulk of it assumes concave dynamics as well as concave utility. Santos (1999) and Santos and Vigo-Aguiar (1998) are excellent entry points into this literature. Our problem appears to not have been treated in the literature because it is a stochastic two dimensional state variable problem with non concave stochastic dynamics. We believe the methods and insights gleaned from our work here may be of general use to economics.

Our model is an instructive example of the tradeoff between material economic production (e.g. agriculture) and other environmental amenities (e.g. clean water for drinking, industry, irrigation, or recreation) (Postel and Carpenter 1997). Various authors (see Brock (1977) and Stokey (1998) and references therein) have maximized the discounted sum of utility

$$U(c,P) = u(c) - v(P), \tag{1}$$

such that

$$\frac{dx}{dt} = -c + f(x,w), \tag{2}$$

$$\frac{dP}{dt} = -mP + w. \tag{3}$$

Here P is pollution stock, w is waste (load), m is natural regeneration, $u(c)$ is utility of material goods, $f(x,w)$ is production of material economic production which can be used for consumption or invested into capital accumulation, x is capital stock (depreciation is embedded into f), and $v(P)$ is disamenity cost of P stock. If we neglect the role of capital, write $f = f(w)$, put $u(c) = \alpha c$, and add a term $r(P)P$ to the P equation, then (3) becomes

$$dP/dt = w + r(P)P - mP. \tag{4}$$

This is a deterministic version of Carpenter, Ludwig and Brock (1999). Therefore we can think of that work as a standard economic growth model with $u(c) = \alpha c$, but with no capital. If we assume that labor L is fully employed but not growing, it may be embedded in $c = f(w) = A(L)w$, with L constant. Likewise we may absorb L into α, so α is the product of average product of a unit of load in production of c and marginal utility of a unit of c in consumption utility.

We also have a slow variable M in the present work. In an economic context it could correspond to gradual degradation of the regenerative capacity, $[r(P, M) - m]$ as M increases. Hence the present work represents a growth model with not only environmental dynamics, but also an additional slow variable that impacts the regeneration rate. The overall environment stock (which is a function of P in our model) represents environmental quality that can be degraded. In the case of water quality, such degradation leads to human health problems, fish kills, loss of recreation and increased costs of treatment before water can be used by municipalities or industries (Postel and Carpenter 1997). "Mud" represents slow wearing down of the ability of the environmental stock to regenerate itself if one keeps loading more P into the system faster than the ability of it to remove at the natural cleansing rate m.

2 Model specification

2.1 Notation

The variables and units are

P - P kg/m^2

L - Load of P kg $m^{-2}yr^{-1}$

M - Mud P kg/m^2

P_c - Half saturation constant for P in (4) $(= 2.4kg/m^2)$

N - noise

r - recycling constant $(= 0.34yr^{-1})$

s - transfer to mud constant $(= 3.3yr^{-1})$

h - flushing constant $(= 0.19 yr^{-1})$

b - removal from mud constant $(= 0.022 yr^{-1})$

$q = 8$.

The values of these parameters used for most computations are indicated in parentheses. The values assigned to certain of these parameters are based upon estimates of Lathrop et al (1998), North Temperate Lakes Long-Term Ecological Research Site (2002), and Reed-Anderson, T., S.R. Carpenter and R.C. Lathrop (2000), but there is considerable uncertainty. We explore some consequences of this uncertainty below, but a full treatment is beyond the scope of this work.

2.2 Dynamics

We describe the dynamics by a difference equation in time, with yearly increments. If rates are large, the difference equations must account for the possibility of several events such as sedimentation and recycling within a single year. Appendix A provides a derivation of these equations from differential equations, as well as simplified versions that hold of the rates s and h are small. The resulting equations are

$$P_{t+1} = e^{-s-h} P_t + \frac{1 - e^{-s-h}}{s+h}[LN_t + rMF(P)], \tag{5}$$

$$M_{t+1} = M_t(1-b) + (1 - e^{-s-h})\frac{s}{s+h}P_t + g(s,h)LN_t + (g(s,h)-1)rM_t f(P_t), \tag{6}$$

$$N_t = \exp(z_t - \frac{1}{2}\sigma^2), \tag{7}$$

$$F(P) = \frac{P^q}{P_c^q + P^q}, \tag{8}$$

where z_t is Normal(0, σ^2) and

$$g(s,h) = [1 - \frac{1 - e^{-s-h}}{s+h}]\frac{s}{s+h} = \frac{s + h - 1 + e^{-s-h}}{s+h}\frac{s}{s+h}. \tag{9}$$

In order to replace the continuous model by a discrete one, we define meshes P_i, $i = 1, \ldots, N_P$ and M_j, $j = 1, \ldots, N_M$.

2.3 Expected discounted present value

We define the expected discounted present value of a policy as the sum

$$v(P,M) = \sum_{t=0}^{\infty} (1-\delta)^t n(t), \tag{10}$$

where $n(t)$ is the expected net return at time t, δ is the discount rate, and P and M are the starting values of the corresponding variables. Given a feedback strategy of loadings that depend upon the current values of these quantities, the expected discounted present value at time t satisfies the dynamic programming equation (Mangel and Clark, 1988):

$$v^t(P_i, M_j) = n(P_i) + (1-\delta)\mathrm{E}[v^{t+1}(P_i^{t+1}, M_j^{t+1})], \tag{11}$$

where P_i^{t+1} and M_j^{t+1} are the new values of P and M, starting at $P = P_i$, $M = M_j$, and E denotes the expectation operator. The dependence of the right-hand side of (**??**) upon the loading $L_{i,j}$ appears in (**??**) of Appendix B. We adopt the same scheme of costs and benefits as in Carpenter, Ludwig and Brock (1999):

$$n(P_i) = \alpha \mathrm{E}[e^z L_{i,j}] - \beta_1 P_i - \beta_2 {P_i}^2. \tag{12}$$

Here α is the benefit per unit of loading, and β_1 and β_2 determine the loss of amenity associated with P dissolved in the water column. We set $\alpha = 1.$, $\beta_1 = 0$, and $\beta_2 = 0.065$ for most cases. It should be noted that current estimates would set $\alpha = 1$, $\beta_1 = 1$, and $\beta_2 = 2.3$ (Stumborg, Baerenklau and Bishop 2001). With such economic parameters, the optimal strategy is nearly trivial: never load the lake at all, except for very low values of P or M. Hence the present work is actually not necessary to manage Lake Mendota optimally according to economic criteria. However our objective is to understand the general problem of lake management, so we have chosen economic parameters in order to produce strategies that might be applicable to other situations. For similar reasons we have also used quite different ecological parameters for the latter part of our results: they illustrate phenomena that are not apparent with the present values. We have chosen to use the actual loadings instead of the target loadings $L_{i,j}$ in (**??**), but other options are possible.

3 Results

Our numerical methods are described in Appendix B. The terms "value iteration" and "policy iteration" are defined there. We carried out the procedures described in Appendix B for a number of choices of parameter values. We used standard parameter values as given in Section 2, with a mesh of 40 points in P and 41 points in M.

3.1 Computation of optimal policies

In each case, the optimal loadings were first computed using the one dimensional model as in Carpenter, Ludwig and Brock (1999). For that case no value iterations are required since the equations for the present value can be solved directly for each value of M. The present value of the one dimensional policy was computed using the two dimensional dynamics. This required some value iterations. After this, policy and value iterations were carried out using the two dimensional model. We found that the value iterations generally converged well, but there were occasional difficulties with convergence of policy iterations, as noted below.

3.2 Description of the optimal policy

The upper left panel of Figure 2 shows that for low values of M (the black curve) the optimal policy is to load heavily regardless of the value of P. An explanation is that the lake can absorb high loadings without appreciable decline in water quality when M is low. By the time that such loadings build up so high that water quality suffers, discounting makes the penalty for poor water quality less important. For a slightly higher value of M, the policy is to use moderate loads unless P is above the threshold value of 2.4. If P is beyond that point, heavy loads are applied. This feature is discussed below in Section 3.8. For higher values of M, the optimal policy shows a steady decrease in loading as P increases. The upper right panel is another depiction of the optimal loadings that allows one to see the qualitative features at a glance. The red band at the bottom shows the high loadings for small M. The feature of increasing loadings as P increases is shown by the increasing thickness of the red band as P increases. For higher values of M the change in colors from greens to blue as P increases shows that the loadings decrease.

The lower left panel of Figure 2 shows some trajectories. They all converge to the black diamond. However, they take a variety of paths to that point. These curves are discussed in more detail below. The lower right hand panel shows the expected discounted present value as a function of P and M. This quantity always decreases as P or M increases.

Figure 3 shows a variety of quantities as function of time for the curves in the lower left hand panel of Figure 2. Note that P comes close to its common equilibrium in the time interval shown, but M does not: it may require over 200 time steps for M to converge. The red curve shows high loading for the first few steps, which results in an initial increase in P. A much bigger increase in P appears in the green curve, but the corresponding loadings are small. This increase is due to recycling from the sediments, since the initial value of P is in the region where recycling is strong and M is large. In contrast, the blue curve shows a decrease in P even though P starts out higher than for the green curve. This is due to low initial M. Note that the green curve is the only one where the present value shows a steady increase. The red and blue curves show an initial decrease in present value, followed by an increase and then a final decrease. This feature would appear more prominently if the scale had not been adjusted to include the green curve.

3.3 Effect of the discount rate on optimal loadings

Figure 4 shows the effect on the optimal loadings of decreasing the discount rate. The upper left hand panel is identical to the corresponding panel in Figure 2. Comparison of the various panels shows a steady decrease in loading as δ decreases: a longer time horizon leads to less loading. Loadings decrease as M increases. These results demonstrate precaution. Another view of this comparison is given in Figure 5.

3.4 Comparison with one dimensional policy

In view of the complexity and estimation difficulties of the two dimensional theory that includes mud dynamics, we might attempt to manage on the basis of the simpler one dimensional theory as in Carpenter, Ludwig and Brock (1999). Figure 6 compares the present values obtained from the two dimensional policy and from the one dimensional policy for $\delta = .01$. There are very substantial differences in loading. Presumably this is because the M dynamics are not negligible: the M dynamics are slower than the P dynamics, but still fast enough that they cannot be neglected. As the lower two panels illustrate, the differences in loadings result in substantial differences in present values.

Figure 7 shows how differences in present value vary as function of the discount rate. The "mean present value" that is plotted is the product of the discount rate and the expected discounted present value. This is a more stable indicator than the present value itself, which varies approximately like the inverse of the discount rate. The differences in present value are substantial, but the variation with the discount rate is comparatively small. This behavior should be contrasted with the behavior for slower dynamics illustrated in Figure 17 below.

3.5 Effect of uncertainty

There are great difficulties in obtaining accurate estimates for the parameters involved in the P and M dynamics. In order to assess the importance of uncertainty in these parameters, we manufactured plausible posterior distributions for each of the parameters r, s, and P_c. We used both a point estimate (called the MLE) and the full posterior to compute loadings. We computed expected present values using the posterior distribution and compared the policies that were derived using the MLE and the posterior. In all cases except for the one involving P_c, the differences in strategy and present value were not large. The posterior in P_c (the approximate location of the threshold in the recycling) includes the values 2.4, 2.8, and 1.99 with equal weights. Figure 8 compares the MLE and posterior optimal loadings and present values for the case $\delta = .05$. The most important difference in loadings appears for $P < 3$, which is the important region for triggering recycling. The abrupt transition from red to dark green for low P indicates a much more precautionary policy than for the MLE optimal policy. These differences result in much larger orange regions in the present value plots below. Figure 9 shows analogous results for $\delta = .01$. The differences in loadings are more pronounced and the differences in present values are proportionally greater.

Figure 10 provides information about variation with the variance of the posterior distribution. We constructed distributions of $\log P_c$ with three points and mean at .8755 (corresponding to $P_c = 2.4$). The variance of the distribution of $\log P_c$ was adjusted to produce a family of posteriors which were used to produce the plots. For low P, the loading decreases as the variance increases, as is shown in the upper panel of Figure 10. Later the loadings increase. This is a very suspicious result. It is associated with a decrease in the present value so sharp that it lies below the value obtained using the MLE. We suspect that this anomaly is due to difficulties in the optimization when there are extreme values of P_c in the posterior: the value of .2 for the standard deviation of $\log P_c$ corresponds to a posterior with P_c values of 2.04, 2.40, and 2.83. The policy iterations fail to converge for this case, unless steps in (**??**) are reduced by using an interpolation between the old value for loading and the new one computed from (**??**). The lower panel shows an initial sharp decrease in present value if uncertainty is neglected.

3.6 Effect of increasing σ^2

The preceding results have all assumed that $\sigma^2 = .01$, which represents a very low variability in actual loadings: see equation (3). Figure 11 shows the effect of varying σ^2. The loadings generally decrease as σ^2 increases, but we show only one case. All of the plots of present values show a very substantial effect of neglecting the variability in loading (σ^2). The present values obtained taking account of variation in loading are all positive, while those that neglect variation in loadings may be very large negative.

3.7 Effect of a linear penalty for P

The optimal loadings shown in Figure 1 all call for high loadings at low levels of M. Is this due to the lack of a penalty term proportional to P? Figure 12 shows the effect of increasing β_1, with $\delta = .05$. Loadings decrease as β_1 increases, as would be expected.

3.8 Effect of reducing rates in the equations

As pointed out in Carpenter, Ludwig and Brock (1999) and in the introduction to this work, the threshold exhibited by the recycling term in (**??**) can lead to two stable equilibria for the dynamics. Under these circumstances the optimal policy may change drastically as a threshold is exceeded. Perhaps the best-known example of such a phenomenon is critical depensation in fisheries (Clark 1976; Clark 1990). Skiba (1978) showed that the optimal policy may show a switch in control in the vicinity of the threshold. Ludwig and Varah (1979) obtained a similar result for a stochastic problem. These results are analogous to the famous demonstration of Clark (1973) that in cases where the rate of return from harvesting a stock sustainably is too low, the optimal economic policy may be to harvest it to extinction. The analogous policy for the present case is to load the lake heavily once P exceeds the threshold level P_c.

In the interest of simplicity, we did not allow strategies that produce two equilibria in Carpenter, Ludwig and Brock (1999), and it is not shown in the results so far. However, if some of the rates are reduced, the behavior changes. Figure 13 shows the effect of changing parameters as follows: $r = .015$, $s = .7$, $h = .15$, $b = .001$, $P_c = 2.4$. The economic parameters are assumed to be $\alpha = 1$, $\beta_1 = 0$, $\beta_2 = .025$, $\delta = .05$. The upper left panel of Figure 13 shows how loadings at intermediate values of M may first decrease, then increase sharply, then decrease again. This is associated with the existence of two stable equilibria, as illustrated in the lower left panel of Figure 13. If P is well below the threshold value of 2.4, it is optimal to keep loadings moderate and hence avoid high values of P. But if M is sufficiently low, then high loadings can be tolerated long enough to outweigh the later losses associated with high P. That is the significance of the high loadings for low P and M. Figure 14 shows additional details of the trajectories in Figure 13.

Figures 15 and 16 present results for $\delta = .01$. The region of high loadings in the upper right panel of Figure 15 is much reduced when compared with Figure 13. In the lower left panel of Figure 15, the trajectories that start at low M now terminate with relatively low P.

Figure 17 is analogous to Figure 7. It compares present values for one dimensional and two dimensional policies for the slow rate parameters. The differences in present values are very large for small discount rates. In contrast to the earlier case with fast dynamics, the differences in present value are insignificant if the discount rate is higher than .015. Evidently there is a strong interaction between the time scale for discounting and the time scale for changes in M when the underlying dynamics are slow. Note that the scale on the horizontal axis is logarithmic. This produces remarkably straight lines in the figure. The details of such figures vary from point to point in the P, M plane: at some points there is very little difference between present values corresponding to the two strategies.

The case of slow rates leads to figures very similar to Figures 10 and 11. Thus random variation in loadings and uncertainty in the value of P_c have effects that are important in both cases.

4 Discussion

4.1 Interaction of economic, ecological and political components

The most striking feature to emerge from the present exercise is the strong interaction between the economic and ecological components of our model. We find drastic differences between the optimal policies and their implications for lake dynamics as economic variables change. This is evident from the comparisons in Figure 4 (where the discount rate varies) and Figure 12 (where the component of loss proportional to P varies). Similar differences appear with the slower dynamics in Figures 13-17. It is difficult to imagine how an adequate account of this system can be given without including the interaction between the economic and ecological components, yet many research programs are not well organized to treat both components together. All of the results cited below depend strongly upon the economic discount rate, as is shown in Figures 4, 5, 7, and 17.

It is evident that public policy is not determined solely by economic considerations, nor should it be. Social effects and political influences are important and cannot be neglected. As in the case of Lake Mendota today, historically based political influence may impede adoption of policies that would otherwise be clearly justified on grounds both of economy and equity (Carpenter et al. 2003).

4.2 The role of slow variables

We previously presented a simpler account of this system that neglected the mud dynamics (Carpenter, Ludwig and Brock, 1999). To what extent are our earlier results valid? The answer depends upon both economic variables and ecological rates. Figures 6 and 7 show a substantial improvement by use of the more complicated model. When using slower rates and low discounting as in Figure 17, there is a substantial effect if discounting is low enough. All of these figures show large differences between the one and two dimensional policies and their present values. In the case of faster rates (Figures 6 and 7) the differences are fairly independent of the discount rate. In the case of slow rates (Figure 17), the difference between policies and present values is important only for low discounting: the mud dynamics are not important unless the time horizon for optimization is long enough for substantial changes to occur. We conclude that the one dimensional policy is adequate only if mud dynamics are slow and the time horizon for optimization is short.

Even the theory developed in this paper may be inadequate for the full P cycle depicted in the upper left panel of Figure 1. A minimal model of soil and fertilization practices was introduced by Carpenter, Brock and Hansen (1999) and is adapted for our purposes as follows. Suppose dynamics of soil P, S, follow

$$S_t + 1 = S_t + (1 - k)F_t - w_t S_t, \tag{13}$$

where F_t is the annual import of P as fertilizer and animal feed in the catchment, k is the proportion of fertilizer and animal feed exported from the basin as farm products each year, and w_t is the fraction of soil phosphorus that runs off to the lake each year.

Runoff occurs in stochastic, weather-driven events. Thus $w_t S_t$ corresponds to loading, and w_t is a random variable chosen to produce the desired distribution of annual loadings. This model ignores inputs of P to the soil from aeolian deposition and weathering. Both fluxes are quite small compared to imports of fertilizer and animal feed (Bennett et al. 1998). In this expanded model, the utility from use of phosphorus will be proportional to kF_t, and the policy control variable will be F_t. Of course, this is a minimal model and a number of additional details could be added.

In the context of our discussion, this model has two interesting features. (1) It adds another slowly-changing variable to the dynamics. Budget studies of soil P in watersheds indicate that the turnover rate is on the order of a century (Bennett et al. 1998, Reed-Anderson et al. 2000). (2) The policy acts through a slow variable, so effects of policy choice have even longer time lags. While analysis of the model with soil dynamics is beyond the scope of this paper, we anticipate that the impacts of discounting will be even greater, due to control through a slow variable.

4.3 The importance of uncertainty and fluctuations in loading

In spite of all the studies of Lake Mendota and similar lakes, we are unable to provide estimates of parameters with any assurance of their accuracy. Is this important? We have done only a preliminary study of uncertainty, but it is already clear that uncertainty in the P_c parameter (that controls the position of the threshold in recycling of phosphorus from sediments) cannot be neglected. Figure 8 (with $\delta = .05$) shows a modest effect of uncertainty, and Figures 9 and 10 show a substantial loss if uncertainty is neglected and a precautionary approach is not adopted. A full treatment would probably reveal higher uncertainty than we have assumed, and consequently a higher cost of neglecting it.

The results for fluctuating loadings in Figure 11 and analogous results for the slower rates show that these fluctuations cannot be neglected. Therefore it appears that a stochastic model is essential.

4.4 Precaution

There has been a great deal of discussion of a "precautionary principle" for public decision-making involving the environment (Bodansky 1991; Dickson 2000; Foster et al 2000). A frequent complaint about such ideas is that they are too vague and they fail to balance costs and benefits. The present results do not suffer from such defects. They are quantitative, although probably not very precise.

A remarkable feature of the present results is that the amount of precaution increases as complications are added to the theory. This is not a new phenomenon. May et al (1978) considered a range of possible forms of stock dynamics which vary in the skewness of the stock-recruit relationship. They evaluated the effect of various levels of exploitation rate upon the distribution of yield, in particular its expectation and coefficient of variation. They found a tradeoff between high yields and variability of yield, and they posed the choice of exploitation rate as a "portfolio" problem. They also called attention to difficulties in stock assessment and the consequent need for management strategies that are robust with respect to uncertainty. We have seen that fluctuations in loading produce substantial reductions in optimal loadings: this result cannot be obtained from a deterministic treatment. Similar phenomena were seen in Ludwig and Varah (1979). Likewise, uncertainty in parameters implies further reductions on loadings. This idea and the related one of occasional catastrophic dynamics were explored in Ludwig (1995) and Ludwig (1998). In most parts of the P, M plane, consideration of the mud dynamics also implies reductions in loading.

It is unfortunate that simple theories always seem to imply overloading. In practical situations, there are sure to be additional complications that we have overlooked. Hence we would advocate that our results not be taken too literally. A large dose of additional precaution should temper public policies.

4.5 Wealth and sustainability

We can use our model with one slow variable and one fast variable to shed some light on economic notions of "sustainable development," "accounting prices," "genuine investment," "natural capital," and the like. See Dasgupta (2001) for the first three notions. Our model can be viewed as a general model of ecosystem management where loading generates benefits for one group of users but generates costs for another group of users. Our model has two state variables: the level of phosphorus in the water (P) and the level of sedimented phosphorus (M). The negatives of these variables can be used to represent two types of "natural capital" in the system. The "resilience" of the faster moving natural capital is compromised by the loss of the slow moving natural capital. This loss of resilience is captured by the increased recycling of P when M is high. Notice that M does not cause any direct loss of utility. It causes indirect loss of utility via its impact on the dynamics of P. We have explored what happens if the manager ignores this feedback from the slow moving variable to the dynamics of P. This exercise could serve as an example of what can happen when the system manager replaces the "true" complex model with a "false" simple model. Figures 7 and 17 help us understand when the losses by this action are large and when they are small.

The optimal expected present value of the managed system at date t given (P_t, M_t) can be written in the form $W(P_t, M_t)$. As we have seen in (11), W satisfies a recursive equation called the "Bellman Equation." If we maximize over the set of all possible loading policies that are functions of the past history of the system it is easy to see that W is a decreasing function of both P and M. This is so because both P and M are "bads" and clearly the maximum value of the system decreases when stocks of bads increase. The decreasing character of W can be seen in Figures 2, 6, 8, 9, 13, and 15. We expect this same property to hold even when the set of policies that we maximize over is restricted. Dasgupta and

Mäler (2001) point out that one can find "accounting prices" at date t of P_t and M_t by differentiating the value function $W(P, M)$ with respect to P and M. These quantities are the discrete time analog of the co-state variables of optimal control theory. These accounting prices are typically nonpositive. They tell the manager how much it is worth to reduce P and/or M at the margin. If we wrote down a continuous time deterministic analog of our system with utility $U = U(L, P)$ and dynamics

$$\frac{dP}{dt} = F_1(P, M) + L, \tag{14}$$

$$\frac{dM}{dt} = F_2(P, M), \tag{15}$$

then Dasgupta and Mäler (2001) and Dasgupta (2001) show how to relate the usual objects of optimal control theory such as the Hamiltonian and the value function, and the Hamilton-Jacobi-Bellman (HJB) equation to notions of "sustainability", "net national product", "social accounting prices", "sustainable measures of income" "value of natural capital" and the like. Here the Hamiltonian is given by $H = \max\{U(L, P) + p_1.(F_1 + L) + p_2.F_2\}$ (where p_1, p_2 are the co-states), the value function is given by $W(P_t, M_t)$, and the Hamilton-Jacobi equation (in current value units) is given by

$$\delta W(P_t, M_t) = \max\{U(L, P) + p_1.(F_1 + L) + p_2.F_2\}, \tag{16}$$

where δ is the discount rate on future utility.

A useful definition of "sustainable policy" would be a policy such that the analog of W does not decline over time (Dasgupta 2001, p. C22). The HJB equation leads automatically to which notion of "income" must correspond to a plausible definition of "sustainable income" i.e. the Hamiltonian becomes a natural notion of "sustainable income." This means, in our context of lake management, that the lake must be managed so that not only is the direct cost of the level of P taken into account, but also the rates of change of P and M must be priced at the accounting prices p_1 and p_2 to correctly account for damages done to the "natural capital" embodied in the lake and its sediment. "Sustainable income" could be defined, for example, as that which can be taken without decreasing the value function of the lake.

Actually one does not need to restrict oneself to optimal policies to develop recursive equations of Bellman type and to develop notions of "accounting prices" like the above (Dasgupta and Mäler (2001)). In fact if one assumes that loadings at date t, L_t are of "feedback form", $L_t = F(P_t, M_t)$, one can develop a Bellman type recursive expression for the expected present value of the system following policy F, call it $W(P_t, M_t; F)$. Similar kinds of expressions can be developed for policies more general than feedback policies. Quantities like W give us a measure of "wealth" (Dasgupta and Mäler, 2001). One can also use measures like W to define a notion of "sustainable management" by defining the management of the lake to be "sustainable" if W does not decrease as time increases. This captures the idea that the power of the lake to produce value is not ever decreasing over time. Dasgupta (2001) shows that this type of management requires $dW/dt = p_1.dP/dt + p_2.dM/dt \geq 0$ for all t. Here $p_1 = \partial W/\partial P$ and $p_2 = \partial W/\partial M$. It will typically be the case that $p_1 < 0$, $p_2 < 0$: it will be so if the policy is optimal. We can see immediately that using a simple model that ignores M and fails to attach the social marginal cost, $-p_2$, to dM/dt is likely to go wrong if sustainable management is the objective. Our computational methods could be used to quantitatively evaluate the effectiveness of different management schemes that use the computed accounting prices p_1, p_2 as inputs to incentive schemes for loaders. The quantitative effect of replacing a complex model that takes into account M-dynamics with a simple model that ignores M-dynamics could be computed for general policies using our recursive methods provided that the policy being evaluated was in recursive form.

Our study has measures of pollutants, e.g. P_t and M_t. But we have not yet discussed a measure of "income". If we view our system as an example of the more general problem of trading off benefits of economic development (letting $u(L)$ the utility of "loading" represent material economic production) against degradation of the ecological service base (letting $v(P)$ represent the cost of environmental degradation) Dasgupta and Mäler (2001) have discussed building an income measure out of the Bellman type expression obtained for the expected present value of following a particular policy.

A measure of "natural capital" for our system could be built by putting $N_{1,t} = P_o - P_t$, $N_{2,t} = M_o - M_t$, where P_o, M_o are "pristine" levels.

Acknowledgments

We thank Kenneth L. Judd for helpful criticism. DL was supported in part by the National Science and Engineering Research Council of Canada under Grant A9239. WAB and SRC are grateful for support from the National Science Foundation. WAB acknowledges support from the Vilas Trust. None of the above are responsible for errors, opinions, or shortcomings of this paper.

Literature Cited

Aghion and Howitt. 1998. Endogenous growth theory, MIT Press: Cambridge, Mass.

Bennett, E.M., T. Reed-Andersen, J.N. Houser, J.R. Gabriel and S.R. Carpenter. 1999. A phosphorus budget for the Lake Mendota watershed. Ecosystems 2: 69-75.

Beyn, W., Pampel, T., Semmler, W. 2000. Dynamic optimization and Skiba sets in economic examples, Department of Economics, University of Bielefeld. http://www.mathematik.uni-bielefeld.de/~beyn/Schwerpunkt/preprints/englisch/preprint_31_00.html

Bodansky, D. 1991. Scientific uncertainty and the precautionary Principle. Environment 33(7): 4-5,44.

Brock, W. 1977. A polluted golden age, in V. L. Smith, ed., Economics of natural and environmental resources, Gordon and Breach: New York, pp. 441-61.

Brock, W., Starrett, D. 2000. Non-convexities in ecological management problems, Discussion Paper, Beijer International Institute of Ecological Economics, also available at http://www.ssc.wisc.edu/~wbrock.

Carpenter, S.R., D. Ludwig and W.A. Brock. 1999. Management of eutrophication for lakes subject to potentially irreversible change. Ecological Applications 9: 751-771.

Carpenter, S.R., W.A. Brock and P.C. Hanson. 1999. Ecological and social dynamics in simple models of ecosystem management. Conservation Ecology 3(2): 4. Available on the internet: http://www.consecol.org/vol3/iss2/art4.

Carpenter, S.R. and M.G. Turner (eds.). 2000. Special Feature on Fast and slow variables in ecosystems. Ecosystems 3: 495-573.

Carpenter, S.R., R.C. Lathrop, P. Nowak, E.M. Bennett, T. Reed-Anderson and P.A. Soranno. 2003. The ongoing experiment: restoration of Lake Mendota and its watershed. Pp. xxx-xxx in J.J. Magnuson and T.K. Kratz (eds.), Lakes in the Landscape. Oxford University Press, London.

Christie, W.J. 1974. Changes in the fish species composition of the Great Lakes. Journal of the Fisheries Research Board of Canada 31: 827-854.

Clark, C. W. 1976. Mathematical bioeconomics: the optimal mangement of renewable resources. John Wiley and Sons: New York.

Clark, C. W. 1990. Mathematical bioeconomics: the optimal mangement of renewable resources, Second edition. John Wiley and Sons: New York.

Clark, C. W. and M. Mangel. 2000. Dynamic state variable models in ecology. Oxford University Press: New York and Oxford.

Dasgupta, P. 2001. Valuing objects and evaluating policies in imperfect economies, Economic Journal, May, pp. C1-C29.

Dasgupta, P., Mäler, K. 2001. Intertemporal welfare economics in imperfect economies, Mimeo, Beijer International Institute of Ecological Economics, Stockholm.

Dechert, W., Brock, W. 2000. Lake game, Department of Economics, The University of Houston and the University of Wisconsin, Madison, Forthcoming in a book edited by , K. Mäler, C. Perrings and D. Starrett.

Deissenberg, C., Feichtinger, G., Semmler, W., Wirl, F. 2000. History dependence and global dynamics in models with multiple equilibria, University of Aix-Marseille II, deissenb@univ.aix.fr; Vienna University of Technology, Bielfeld and New School University, University of Vienna.

Dent, C.L., G.S. Cumming and S.R. Carpenter. 2002. Multiple states in river and lake ecosystems. Philosophical Transactions of the Royal Society, Series B. in press.

Dickson, B. 2000. Precaution at the heart of CITES? pp. 38-46 in J. Hutton and B. Dickson eds. Endangered species threatened convention: the past, present and future of CITES. London: Earthscan.

Foster, K. R., P. Vecchia and M. H. Repacholi. 2000. Science and the precautionary principle. Science 288 (12 May 2000): 979-981.

Golley, F.B. 1993. A History of the ecosystem concept in ecology: More than the sum of the parts. Yale University Press, New Haven CT.

Holling, C. S. 1973. Resilience and stability of ecological systems. Ann. Rev. Ecol. Syst. 4: 1-23.

Lathrop, R.C., S.R. Carpenter, C.A. Stow, P.A. Soranno and J.C. Panuska. 1998. Phosphorus loading reductions needed to control blue-green algal blooms in Lake Mendota. Canadian Journal of Fisheries and Aquatic Sciences 55: 1169-1178.

Levin, S.A. 1992. The problem of pattern and scale in ecology. Ecology 73: 1943-1967.

Ludwig, D. 1995. A theory of systainable harvesting. SIAM J Appl. Math. 55(2): 564-575.

Ludwig, D. 1998. Management of stocks that may collapse. Oikos 83: 397-402.

Ludwig, D. and J. M. Varah. 1979. Optimal harvesting of a randomly fluctuation resource II: Numerical methods and results. SIAM J. Appl. Math. 37: 185-205.

Mangel, M. and C. W. Clark. 1988. Dynamic modeling in behavioral ecology. Princeton University Press: Princeton, New Jersey.

May, R. J. R. Beddington, J. W. Horwood and J. G. Shepherd. 1978. Exploiting natural populations in an uncertain world. Math. Biosci. 42: 219-252.

Moran, M. and J. Maroto. 2001. Non-smooth dynamic optimization, Department of Economic Analysis, Madrid, Spain.

North Temperate Lakes Long-Term Ecological Research Site. 2002. Publicly-available data sets. http://lter.limnology.wisc.edu.

O'Neill, R.V., D.L. DeAngelis, J.B. Waide and T.F.H. Allen. 1986. A Hierarchical concept of ecosystems. Princeton University Press, Princeton NJ.

Postel, S. and S.R. Carpenter. 1997. Freshwater ecosystem services. pp. 195-214 in G. Daily (ed.), Nature's Services. Island Press, Washington D.C.

Press, W. H. et al. 1986. Numerical recipes. Cambridge University Press: Cambridge and New York.

Reed-Anderson, T. 1999. Flows of phosphorus on an agricultural landscape: Implications for eutrophication and restoration of Lake Mendota. Ph.D. thesis, University of Wisconsin-Madison, USA.

Reed-Anderson, T., S.R. Carpenter and R.C. Lathrop. 2000. Phosphorus flow in a watershed-lake ecosystem. Ecosystems 3: 561-573.

Santos, M. S. 1999. Numerical Solution of Dynamic Economic Models. in J. B. Taylor and M. Woodford, eds., Handbook of macroeconomics, Elsevier Science, B. V. Amsterdam, pp. 311- 386.

Santos, M., Vigo-Aguiar, J. 1998. Analysis of a numerical dynamic programming algorithm applied to economic models, Econometrica 66(2): 409-426.

Schindler, D.W. 1977. The evolution of phosphorus limitation in lakes. Science 195: 260-262.

Semmler, W., Sieveking, M. 1999. Using vector field analysis for studying debt dynamics, Department of Economics, University of Bielefeld.

Skiba, A. 1978. Optimal growth path with a convex-concave production function. Econometrica 46: 527-540.

Stokey, N. 1998. Are there limits to growth? International Economic Review. 39(1): 1-30.

Stommel, H. 1963. Varieties of oceanographic experience. Science 139: 572-576.

Stumborg, B. E., K. A. Baerenklau and R. C. Bishop. 2001. Nonpoint source pollution and present values: a contingent valuation of Lake Mendota. Review of Agricultural Economics, 23: 120-132.

Turner, M. G., R. H. Gardner and R. V. O'Neill. 2001. Landscape ecology in theory and practice: pattern and process. Springer-Verlag, NY.

Weitzman, M. 2002. Income, capital, and the maximum principle, book length manuscript, Department of Economics, Harvard University.

Appendices

A Derivation of the difference equations

If the rates of sedimention (s) and flushing (h) are not small, then there may be competition between the corresponding processes: the simple model that is linear in these parameters is not appropriate. For instance, if

$$P^{t+1} = P^t + L^t N^t - (s+h)P^t + rM^t F(P^t), \tag{17}$$

$$M^{t+1} = M^t + sP^t - bM^t - rM^t F(P_t), \tag{18}$$

and $s+h > 1$, then P may become negative, which is impossible. Appendix A of Carpenter, Ludwig and Brock (1999) addressed this issue by deriving the difference equation from a differential equation. The analogous equation for P would be

$$\frac{dP}{dt} = L - (s+h)P + rMF(P). \tag{19}$$

If we ignore changes in L and the recycling term (this is by no means justified except by expediency), the integrating factor is $\exp[(s+h)t]$. This yields

$$P_{t+1} = e^{-s-h}P_t + \frac{1-e^{-s-h}}{s+h}[L + rMF(P)]. \tag{20}$$

Notice that if $s+h$ is small, then (**??**) reduces to (**??**). The loss in P due to sedimentation and flushing is $(1-e^{-s-h})P_t$. The fraction of this due to sedimention (namely $s/(s+h)$) adds to the M equation. A similar argument applies to the terms involving loading and recycling. Hence

$$M_{t+1} = M_t(1-b) + (1-e^{-s-h})\frac{s}{s+h}P_t + g(s,h)L + (g(s,h)-1)rM_t f(P_t), \tag{21}$$

where

$$g(s,h) = [1 - \frac{1-e^{-s-h}}{s+h}]\frac{s}{s+h} = \frac{s+h-1+e^{-s-h}}{s+h}\frac{s}{s+h}. \tag{22}$$

Note that if s and h are small, then $g(s,h) \approx s/2$. Hence (**??**) reduces to (**??**) if s and h are small, except for a term $Ls/2$, which represents the load that is sedimented during the year.

B Algorithm for calculation of the expected discounted present value

We obtain the optimal feedback and the associated present value by solving the Bellman equation of stochastic dynamic programming (Mangel and Clark, 1988; Clark and Mangel, 2000). If we replace the continuous variables P and M by sets of values defined on meshes, then the continuous dynamic model becomes a discrete one that is amenable to computation. In contrast to Mangel and Clark, we seek solutions that are independent of time. These solutions may be obtained as limits of the usual time-dependent solutions as the time horizon recedes to infinity.

We solve the Bellman equation by a combination of so-called value iterations and policy iterations. We begin with an initial guess for the policy and present value that is analogous to the results in Carpenter, Ludwig and Brock (1999). Value iteration uses the initial guess for present value as the final payoff after T years, and the value at earlier years is obtained by the usual backwards iteration of dynamic programming, using a fixed policy. As $T \to \infty$, the value at the initial time approaches a steady state, which is the result of the value iteration. This steady state is the basis of a policy iteration: a new policy is determined to maximize the value obtained by the preceding value iteration. The value of this new policy is then computed by value iteration, and then the policy is updated by successive applications of this process until there are no significant changes in policy or value.

Although such a procedure might seem cumbersome when compared with simulation results, it should be kept in mind that comparable results using simulations would require many simulations (perhaps hundreds) corresponding to each combination of P and M. We have used simulations to check the present method.

B.1 Approximation of the Expectation of future Values

We approximate the function v^{t+1} in (**??**) by a linear interpolation of its values at mesh points. Since z in (**??**) is unbounded, values of v^{t+1} beyond the last mesh point also enter in (**??**). We extrapolate v^{t+1} to decay quadratically beyond that point, proportional to $\beta_2 P^2$. This ensures that P concentrations higher than the largest mesh point are penalized. Let

$$P_t^* = e^{-s-h} P_t + \frac{1-e^{-s-h}}{s+h} r M^t F(P^t). \tag{23}$$

This quantity may be interpreted as the P level that would be expected in the absence of loading. Since loading is non-negative, all accessible mesh points must lie at or above P_t^*. Let z_k correspond to the random factor required to bring P^{t+1} to the k-th mesh point. If $P_k > P_t^*$, then (**??**) implies that

$$z_k = \frac{1}{\sigma} \log\left(\frac{P_k - P_t^*)}{L_{i,j}^*}\right), \tag{24}$$

where

$$L_{i,j}^* = L_{i_j} \frac{1-e^{-s-h}}{s+h}. \tag{25}$$

It is convenient to extend the definition of z_k to cover the first interval that contains P_t^* and the interval beyond the last mesh point: let

$$z_{k^*} = -\infty \qquad \text{if } P_{k^*} \leq P_t^* < P_{k^*+1}, \tag{26}$$

$$z_{N+1} = \infty. \tag{27}$$

B.1.1 Linear interpolation

If P^{t+1} is between P_k and P_{k+1}, then the linear interpolation of v^{t+1} (denoted by v henceforth) is

$$v(P^{t+1}) \approx \frac{v_k P_{k+1} - v_{k+1} P_k + P^{t+1}(v_{k+1} - v_k)}{P_{k+1} - P_k}. \tag{28}$$

From (**??**) we have

$$P^{t+1} = P_t^* + L_{i,j}^* e^z. \tag{29}$$

Hence

$$v(P^{t+1}) \approx A_k + B_k e^z + C_k e^{2z}, \quad \text{for} \quad k^* \leq k \leq N, \tag{30}$$

where

$$A_k = \frac{v_k P_{k+1} - v_{k+1} P_k + P_t^*(v_{k+1} - v_k)}{P_{k+1} - P_k}, \tag{31}$$

$$B_k = L_{i,j}^* \frac{v_{k+1} - v_k}{P_{k+1} - P_k}, \tag{32}$$

$$C_k = 0, \tag{33}$$

provided that $k^* \leq k < N$. The definition of C_k is for later convenience.

B.1.2 Extrapolation

We extrapolate v for $P^{t+1} > P_N$ as follows :

$$v(P^{t+1}) = v_N - \frac{\beta_2}{\delta}\left[(P^{t+1})^2 - P_N^2\right]. \tag{34}$$

This leads to

$$v(P^{t+1}) = A_N + B_N e^z + C_N e^{2z}, \tag{35}$$

where

$$A_N = v_n - \frac{\beta_2}{\delta}\left[(P_t^*)^2 - P_N^2\right], \tag{36}$$

$$B_N = -\frac{2\beta_2}{\delta} L_{i,j}^* P_t^*, \tag{37}$$

$$C_N = -\frac{\beta_2}{\delta} {L_{i,j}^*}^2. \tag{38}$$

With these definitions and approximations, it follows that

$$\mathrm{E}[v(P^{t+1})] \approx \sum_{k=k^*}^{N} \frac{1}{\sqrt{2\pi}} \int_{z_k}^{z_{k+1}} \exp(-\frac{1}{2}z^2)\left(A_k + B_k e^z + C_k e^{2z}\right) dz. \tag{39}$$

The integrals may be evaluated in terms of the cumulative distribution function for the normal distribution:

$$\Phi(u) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{u} \exp(-\frac{1}{2}z^2) dz. \tag{40}$$

This yields

$$\begin{aligned}\mathrm{E}[v(P^{t+1})] \approx \sum_{k=k^*}^{N} & A_k \left[\Phi(z_{k+1}) - \Phi(z_k)\right] \\ & + B_k \exp(\frac{1}{2}\sigma^2) \left[\Phi(z_{k+1} - \sigma) - \Phi(z_k - \sigma)\right] \\ & + C_k \exp(2\sigma^2) \left[\Phi(z_{k+1} - 2\sigma) - \Phi(z_k - 2\sigma)\right].\end{aligned} \tag{41}$$

These equations determine v^t for a specified $L_{i,j}$ and v^{t+1}.

B.1.3 Interpolation in M

At this point we must face the fact that M_ℓ actually depends upon (P_i, M_j) according to (**??**). In general M_ℓ lies between two mesh points; ℓ is not an integer, but $v^{t+1}(P_k, M_\ell)$ can be obtained by linear interpolation between mesh points.

B.2 Value (backwards) iteration

If the value v^T is specified at a final time T, then (**??**) may be solved successively to obtain $v^{T-1}, \dots, v^0$. If $T \to \infty$ with the end value v^T fixed, then discounting ensures that the sequence of the corresponding v^0 converges. This limit is the solution of the time-independent Bellman equation obtained from (**??**) by deleting the superscripts.

B.3 Consideration of the posterior distribution

In the case where parameter values are given by a posterior distribution, the preceding equations must be modified. We perform value iterations for each point of the posterior: the value function is then given by

$$v(P_i, M_j) = \sum_p w_p v^p(P_i, M_j), \tag{42}$$

where v^p denotes the limit of the preceding value iterations for a single point p in the posterior and w_p is the weight associated with the p-th point of the posterior.

B.4 Policy Iteration

So far the feedback policy $L_{i,j}$ was arbitrarily prescribed. We can use the Bellman equation to improve a given policy. Given a policy L^τ, and the corresponding value $v^{t,p,\tau}(P_k, M_l)$, we can apply (**??**) to each term to obtain

$$v^{t+1}(P_i, M_j, L) = \sum_p w_p \left(n^p(P_i, L) + (1-\delta)\mathrm{E}\left[v^{t,p,\tau}(P_k(L), M_l) \right] \right). \tag{43}$$

Here L stands for a matrix of target loadings. Policy iterations consist in starting with a guess L^0 for the loading and calculating the associated value $v^{t+1}(P_i, M_j, L)$. We define the next approximate loading by maximizing the right-hand side of (**??**). Thus the next loading $L^{\tau+1}$ satisfies

$$L_{i,j}^{\tau+1} = \arg\max_L \sum_p w_p \left(n^p(P_i, L) + (1-\delta)\mathrm{E}\left[v^{t,p,\tau}(P_k(L), M_l) \right] \right). \tag{44}$$

This maximization is carried out separately for each point P_i, M_j. Note that the complete sum over the posterior is required for each maximization. We used Brent's method, first evaluating the function on a mesh of 50 points in order to try to find all local maxima. The simpler method of looking for a single local maximum fails because the right-hand side of (**??**) may have several local maxima.

C Kuznets dynamics

There is evidence in, for example, cross country data, that there is an inverted U relationship (called an "environmental Kuznets curve") between various measures of pollution and per capita income (Stokey (1998)). The idea is that when a country is poor it is willing to trade off its initially clean environment for economic progress. But once it becomes wealthier, it begins to care more about cleaning up its degraded environment than having more material things. The studies cited by Stokey seem to have documented a fairly consistent inverted U shaped relationship between measures of several different pollutants and a measure of income. A standard intuition behind these patterns is as follows. When a country is poor the marginal utility of material consumption is very high relative to the marginal utility of consumption of environmental services. To put it another way there is initially a lot of environmental services but very little in the way of material comforts. Hence the country uses dirty technologies that are more productive than clean technologies to produce material things and accumulate productive power until the marginal utility of consumption of material things gets more in line with the marginal utility of consumption of environmental services. Then the country switches to cleaner technologies (that are less productive), the environment regenerates itself and/or the society invests some of the now plentiful material productive power into cleaning up their environment. Thus an inverted U-shaped pattern emerges between pollution and income for appropriate parameter values. See Stokey (1998) and Weitzman's new textbook (Weitzman (2002)) for worked out examples in an optimal control context.

This raises the issue of whether there is an analog of such patterns in our current study. To put it another way, is there a set of parameter values for our system such that loadings are initially large for small initial levels of P and M so that P, in particular, increases from a low initial level, takes a peak at some time, and then falls as time continues to increase? Our utility of L is linear in L. Hence the analog of marginal utility of consumption in our case is constant, not high with low loading and low with high loading. Thus an important cause of an inverted U pattern is missing in our model. However, if one starts from a degraded state of the lake, it is possible to find parameter values and initial levels of P and M where the value W rises initially, and then falls over a segment of time. This is prominent in Figure 3, where the green curve shows a long increase in W followed by a gradual decrease.

We remark that even though we assumed utility is linear in L in this paper, it would be fairly straightforward, given the work we have already done, to study a generalization where marginal utility of L is falling. We could also generalize the model to allow mitigation and cleanup technologies. These generalizations are deferred to future research.

The pattern of initial rise in P from a low level to a peak level, followed by a fall can be caused by high loading. This is shown in the red curves in Figures 3, 14, and 15. When there is less discounting, as in Figure 15, P does not rise so high. This seems to be more due to an "overshoot" due to discounting in our model than due to a large initial desire for the benefits of loading when "poor" in any sense. An "inverted U" temporal pattern of P arises when initial mud is low so that there is available a period of loading before mud rises, but when it finally does rise and start to cause an increase in P, discounting makes this future problem a small cost in current value dollars. If the initial level mud is already high enough, this period of relatively cheap loading is not available and we do not observe the "overshoot." The presence of the slow variable, mud, appears critical to generate this kind of temporal pattern.

There is a different cause of some of the high P values produced: if P is initially above the threshold for recycling and M is initially high enough, then P can rise even under conditions of no loading. This is shown in the green curve in Figure 3, and the blue curves in Figures 14 and 16.

D Figures

Figure 1:

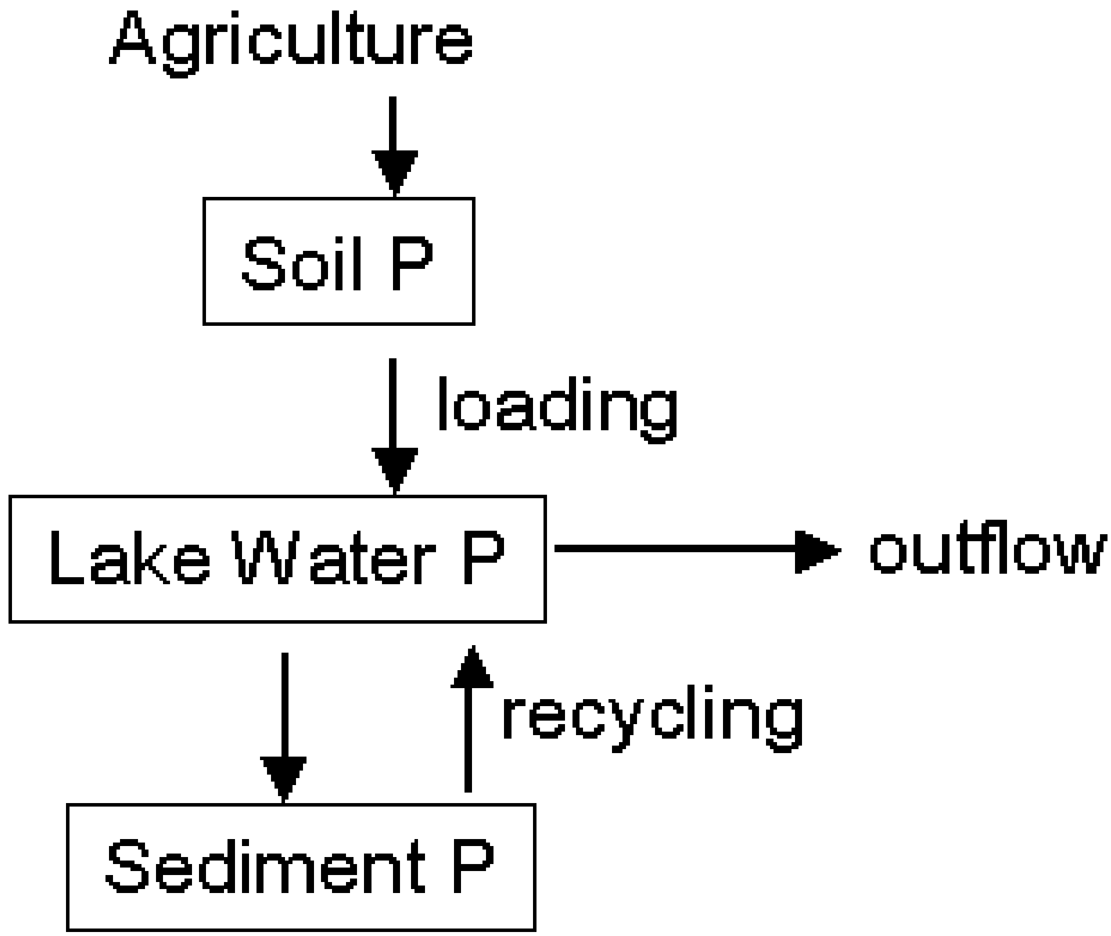

Figure 2: The upper left panel shows optimal loading as function of P for fixed M, for economic discount rate $\delta = .05$. The curves colored black, red, green, and blue correspond to $M = 0$, 5, 15, and 45, respectively. The upper right hand panel displays the optimal loadings as functions of P and M with values corresponding to colors in the key on the far right. The lower left hand panel shows some trajectories in the P, M plane if optimal loadings are applied. The colors of these trajectories refer to curves in Figure 3. All of the trajectories terminate in the point marked with a black diamond. The lower right panel shows the expected discounted present value as function of P and M. The region with negative present values is not colored.

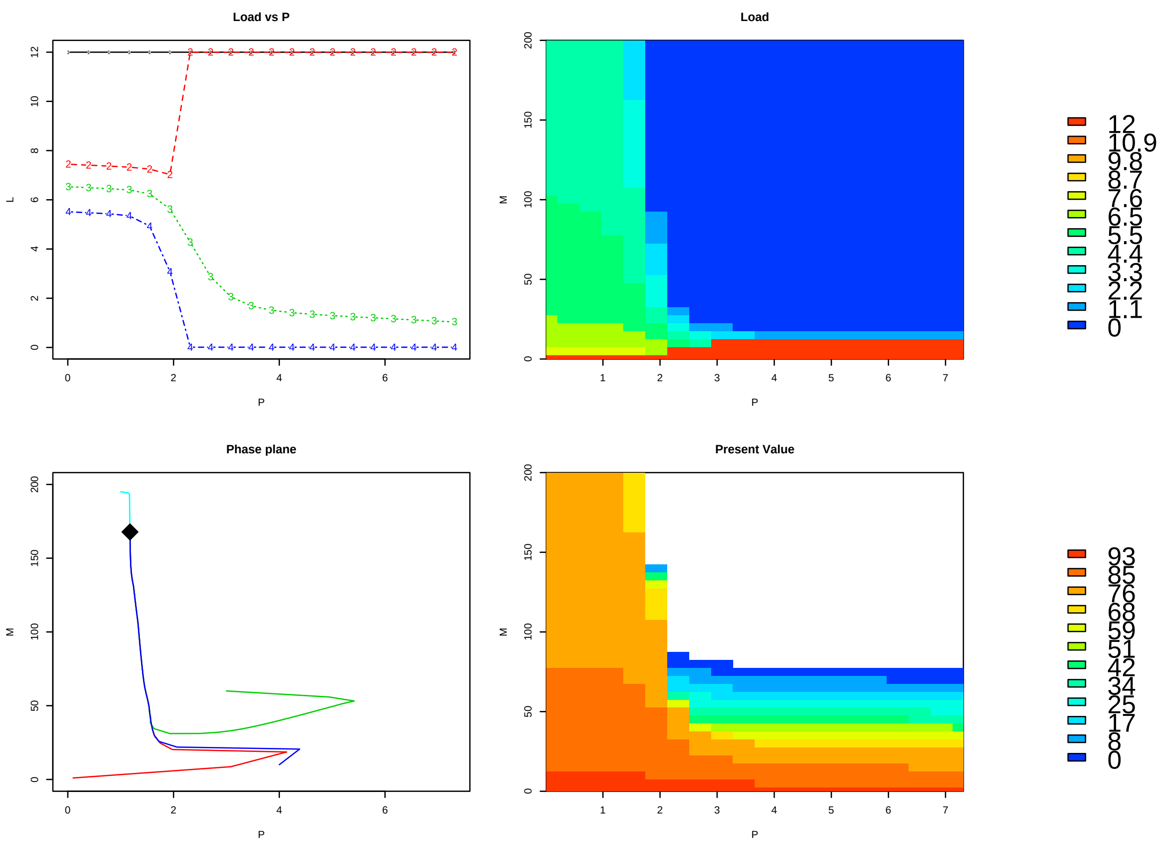

Figure 3: P, M, Load, and Present Value as functions of time, for the trajectories shown in the lower left panel of Figure **??**.

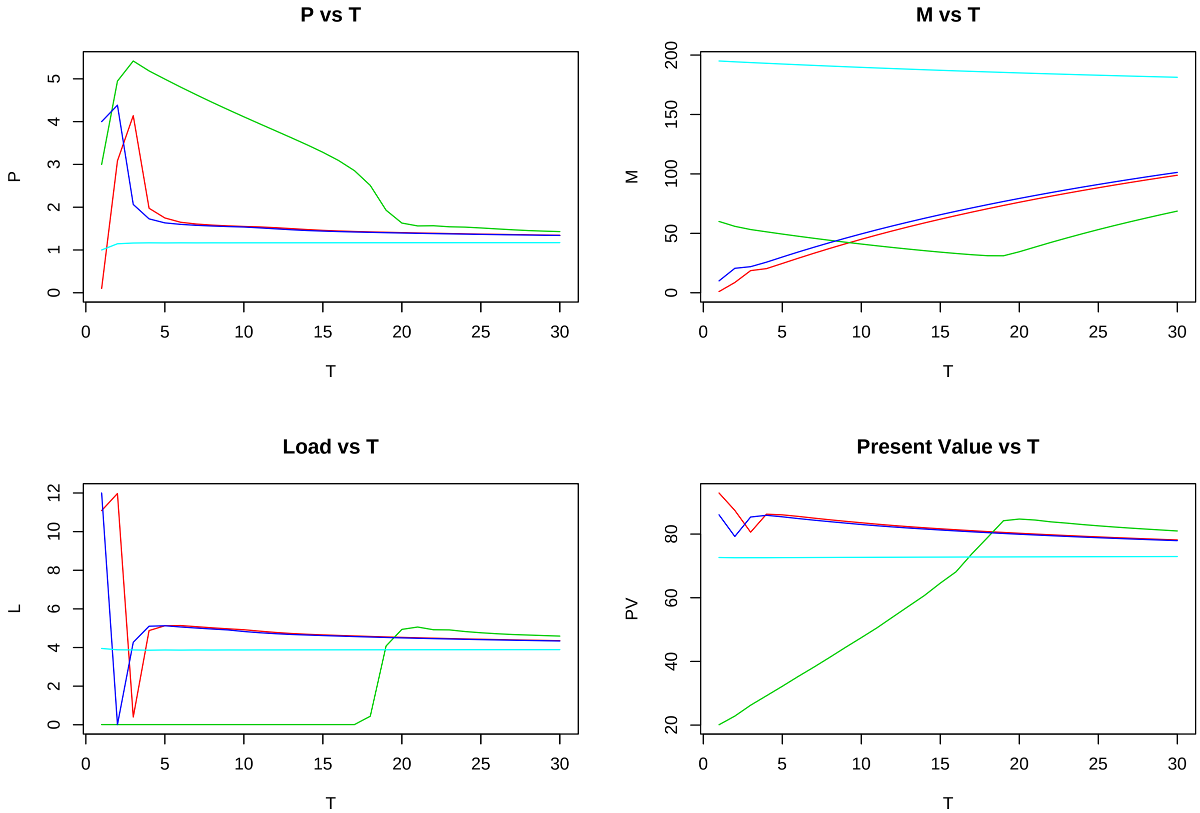

Figure 4: Optimal loadings as function of P for various M and δ. The respective values of M are displayed in the legend in the lower right hand panel, and the values of δ are in the titles of the panels.

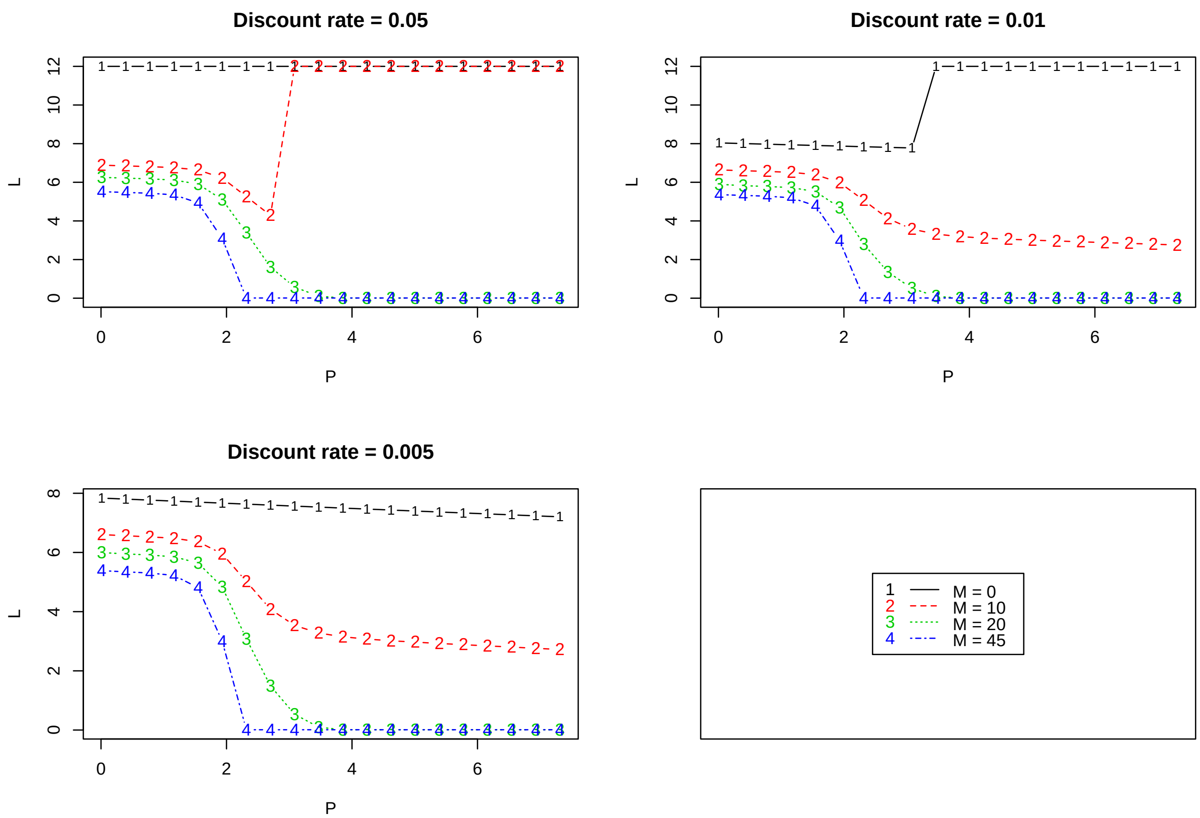

Figure 5: Loadings as function of P and M, for $\delta = .05$, .01,, and .005.

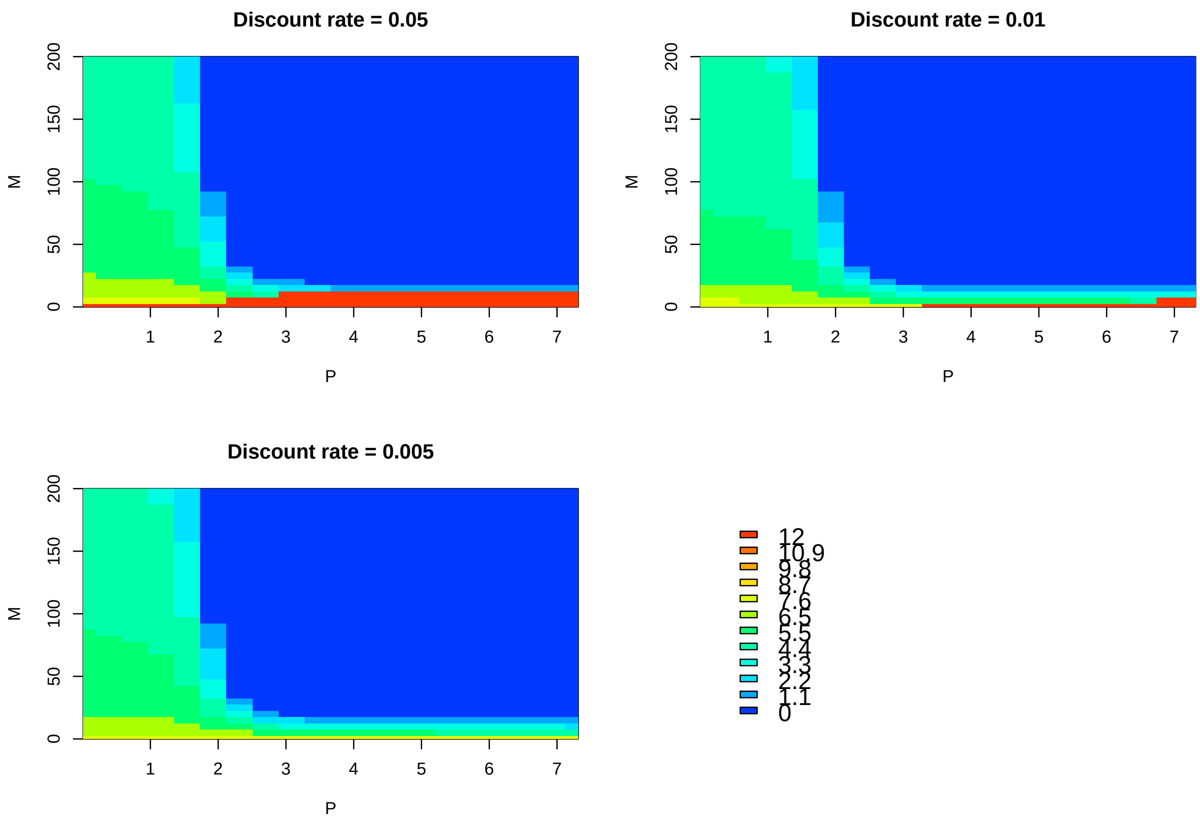

Figure 6: Comparison of one and two dimensional optimal loadings and present values for discount rate $\delta = .01$. The upper left panel shows the optimal load for the two dimensional model as function of P and M; the upper right hand panel shows the optimal load for the one dimensional model. The key on the extreme right shows the values corresponding to the colors. The lower panels show the corresponding expected discounted present values.

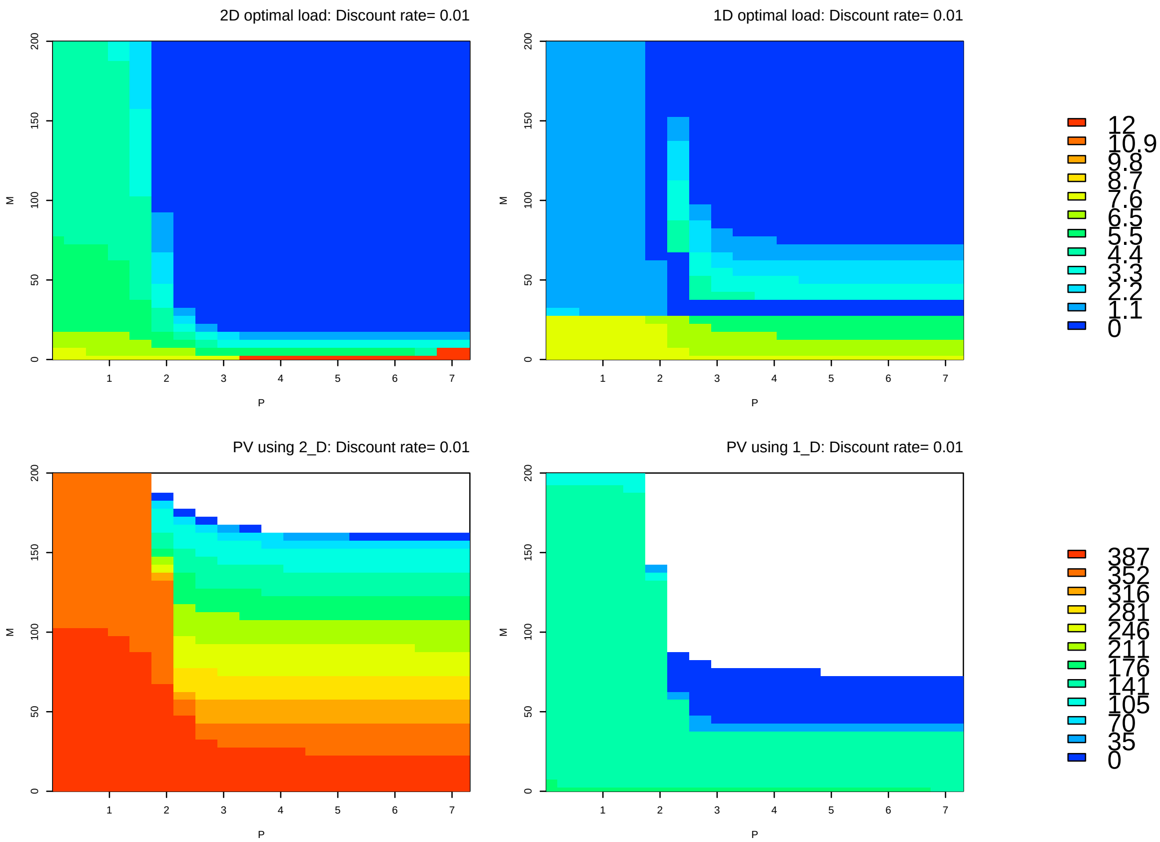

Figure 7: The upper panel shows the mean returns for $M = 2, P = 0$ as computed using the two dimensional policy (solid line) and the one dimensional policy (dotted line) as functions of the discount rate. The lower panel displays the difference in these mean returns. The scale on the horizontal axis is logarithmic.

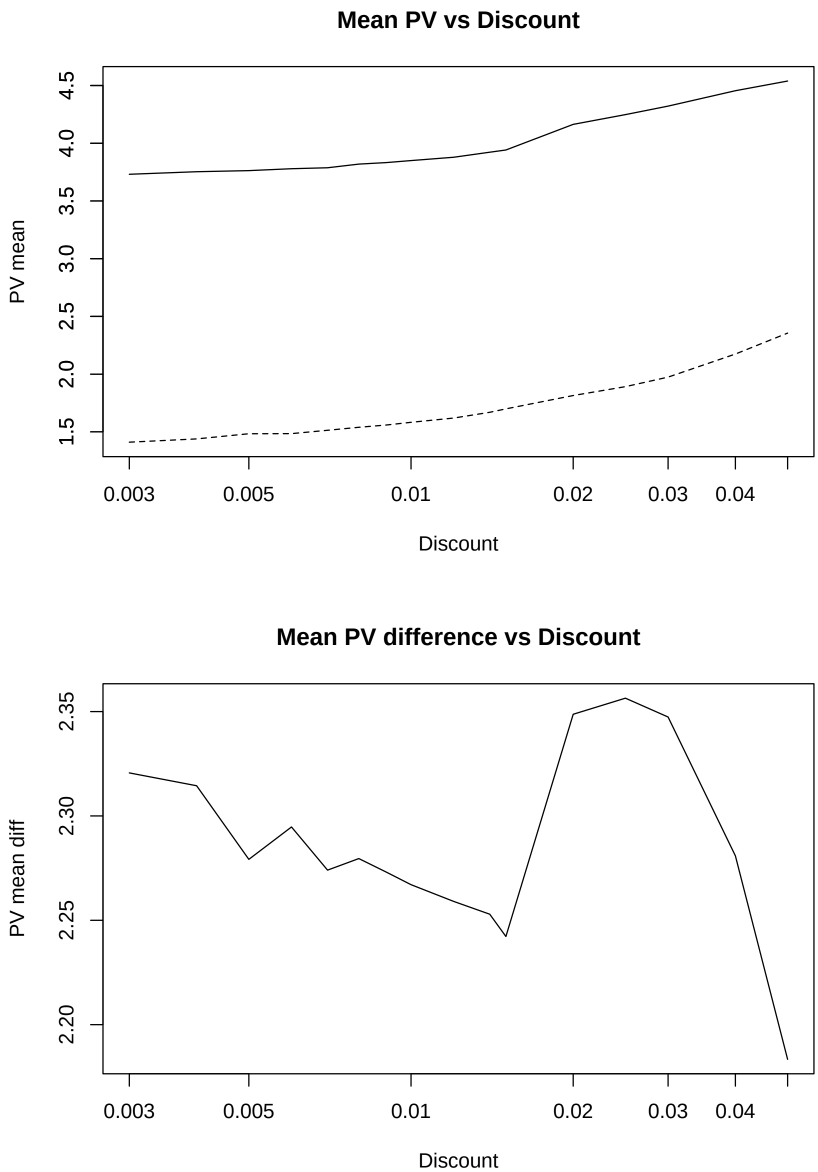

Figure 8: Comparison of optimal loadings and present values for policies taking account of and neglecting uncertainty in P_c, for $\delta = .05$.

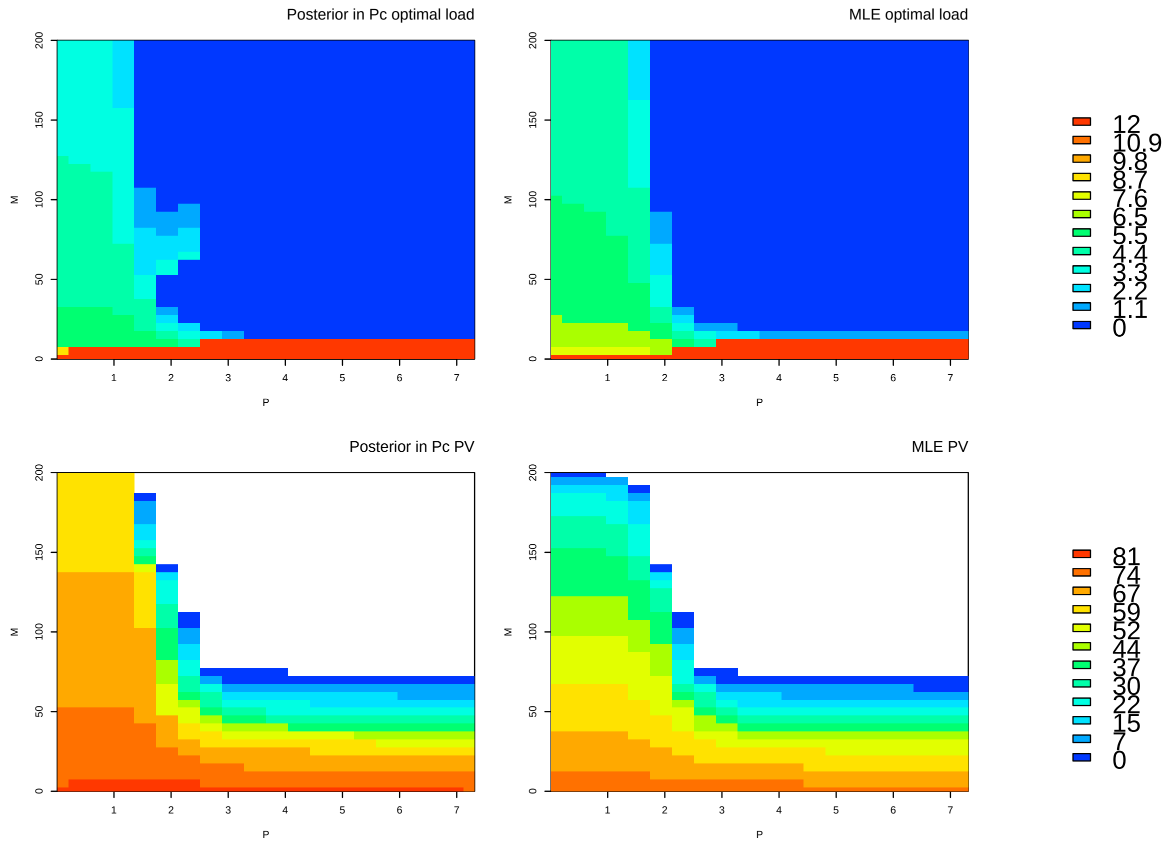

Figure 9: Comparison of optimal loadings and present values for policies taking account of and neglecting uncertainty in P_c, for $\delta = .01$.

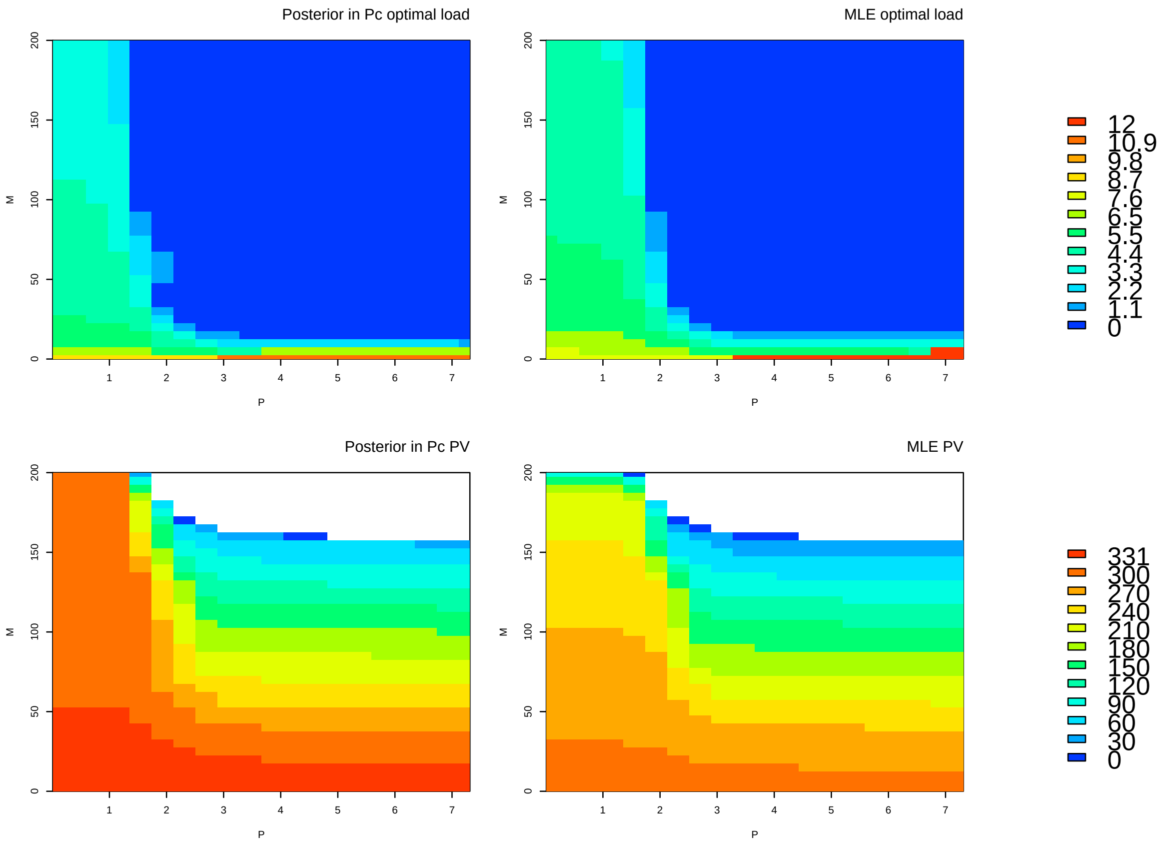

Figure 10: The upper panel shows the optimal loading for $M = 5$, $P = 0$ as function of the variance in the distribution of $\log P_c$, while the lower panel shows the expected discounted present value at that point as function of the loading. The dotted curve in the lower panel shows the expected discounted present value (computed using the posterior of $\log P_c$) if the loading is optimal according to a point estimate.

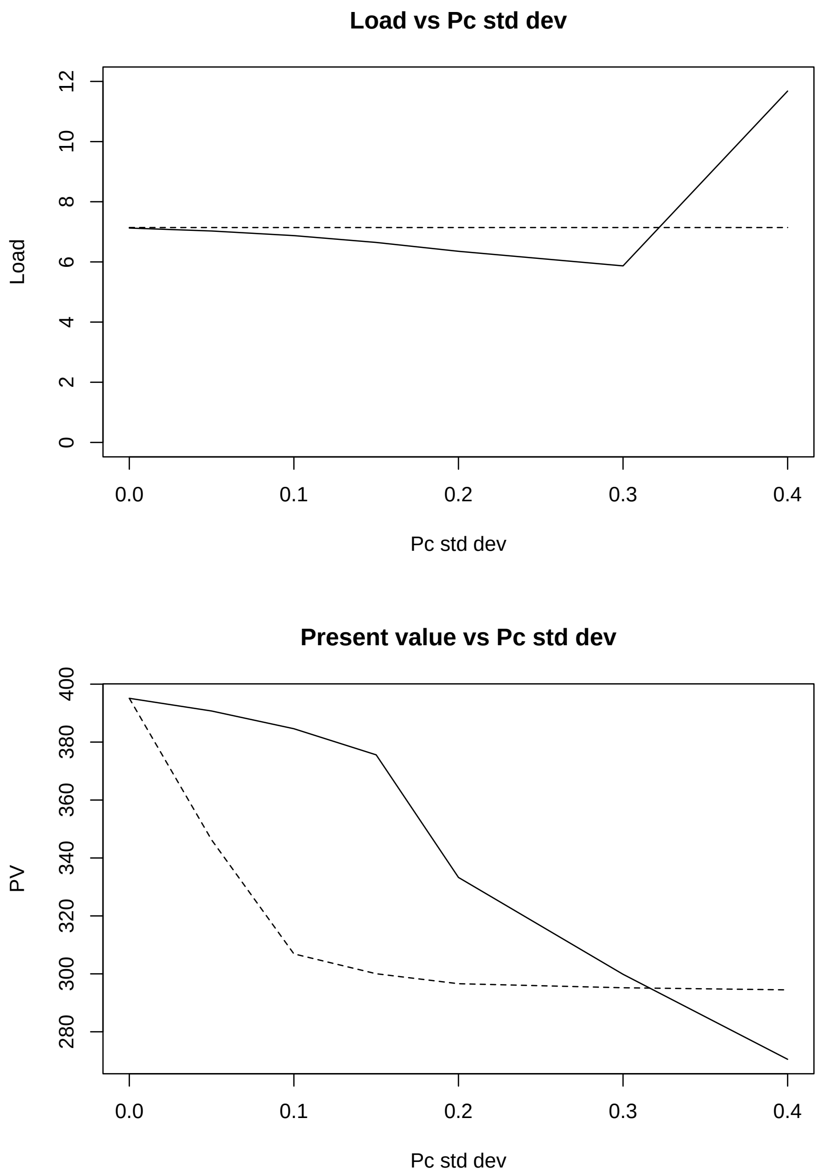

Figure 11: The upper panel shows the optimal loading for $M = 5$, $P = 0$ as function of the variance in the loading σ^2, while the lower panel shows the expected discounted present value as function of the variance in the loading. The dotted curve in the lower panel shows the expected discounted present value if the loading is optimal for $\sigma^2 = 0$.

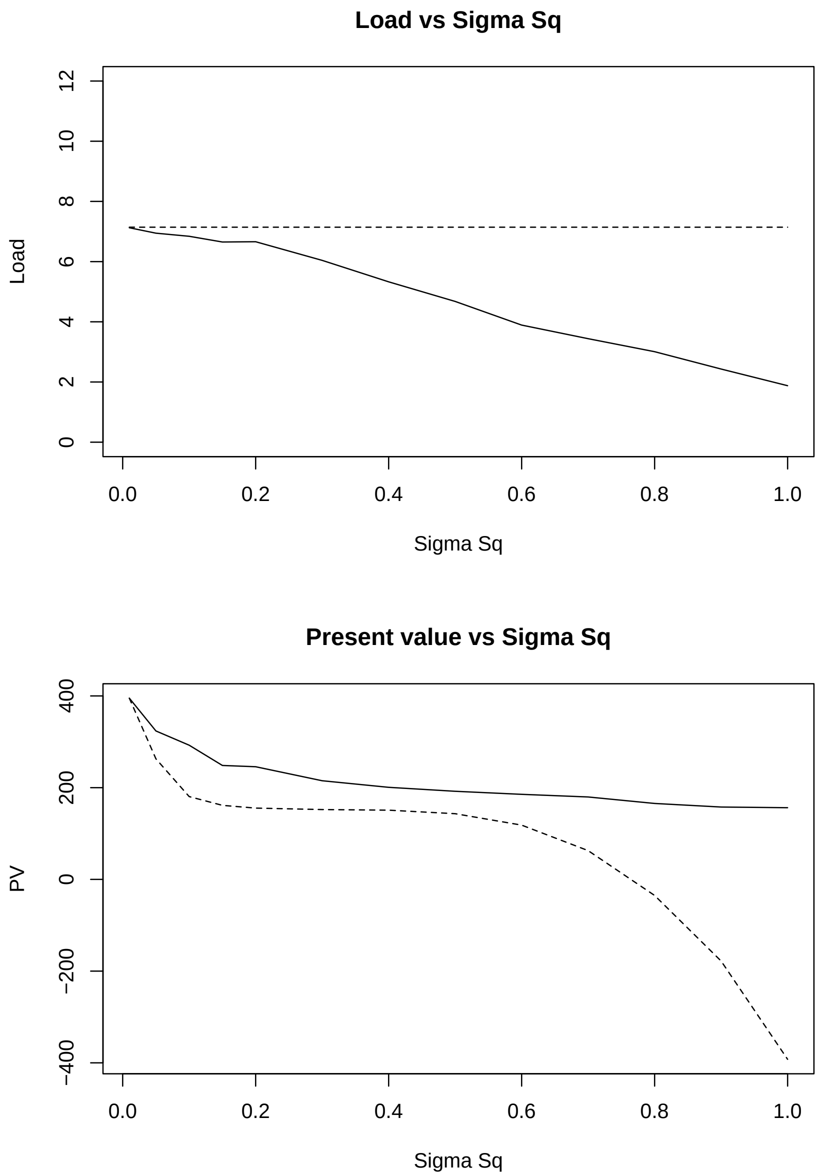

Figure 12: Optimal loading as function of P and M for $\delta = .05$. The values of β_1 used are shown in the labels above each plot

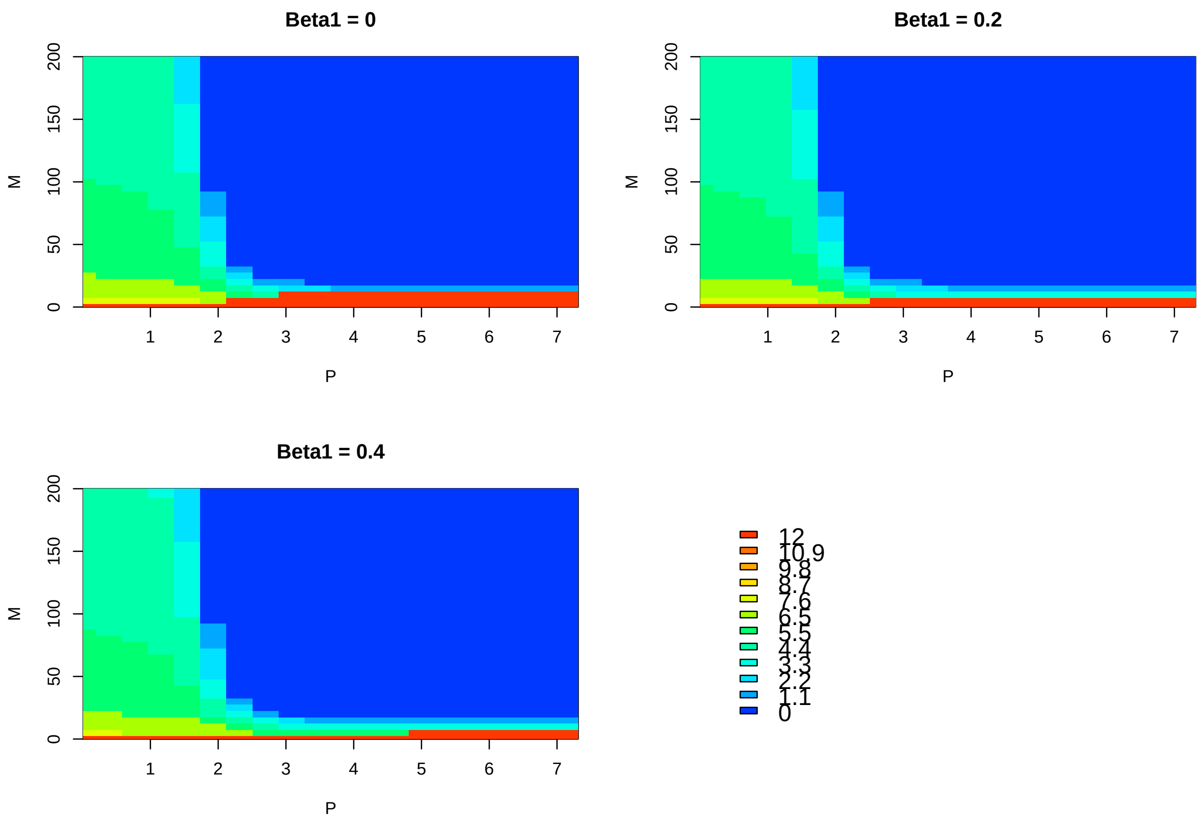

Figure 13: This figure is analogous to Figure 2, but it uses lower rates, as explained in Section 3.8. The black, red, green, and blue lines in the upper left panel correspond to $M = 0$, 360, 720, and 1080, respectively. In the lower left panel, the blue and red areas indicate different domains of attraction. The respective limits are marked with diamonds. Four trajectories are shown, all colored white.

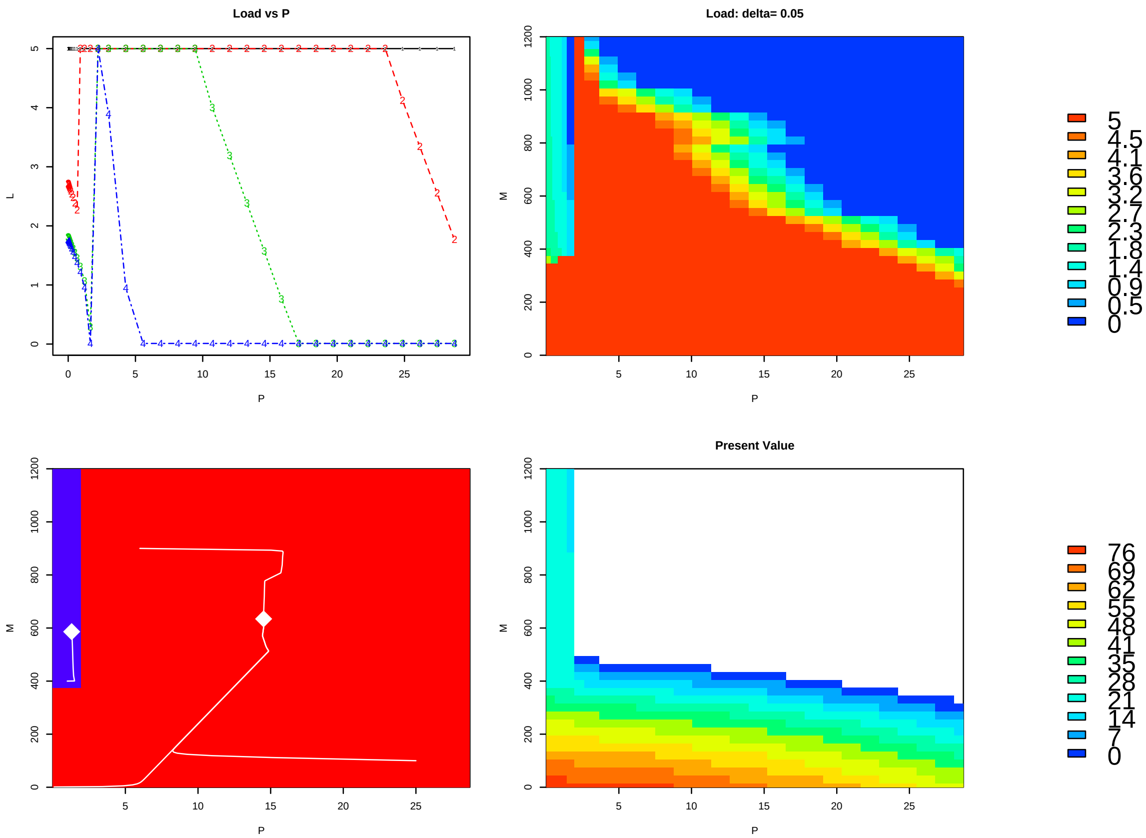

Figure 14: The trajectories correspond to those in the lower left panel of Figure 13. The dark blue trajectory starts in the upper part of the red area in Figure 13, and the red trajectory starts in the lower left part of the red area. The green trajectory starts in the lower right part of the red region in Figure 13, and the cyan curve starts in the lower part of the blue region. The time interval used here is too short to show that the cyan trajectory eventually reaches the diamond in the blue region and the other trajectories reach the diamond in the red area.

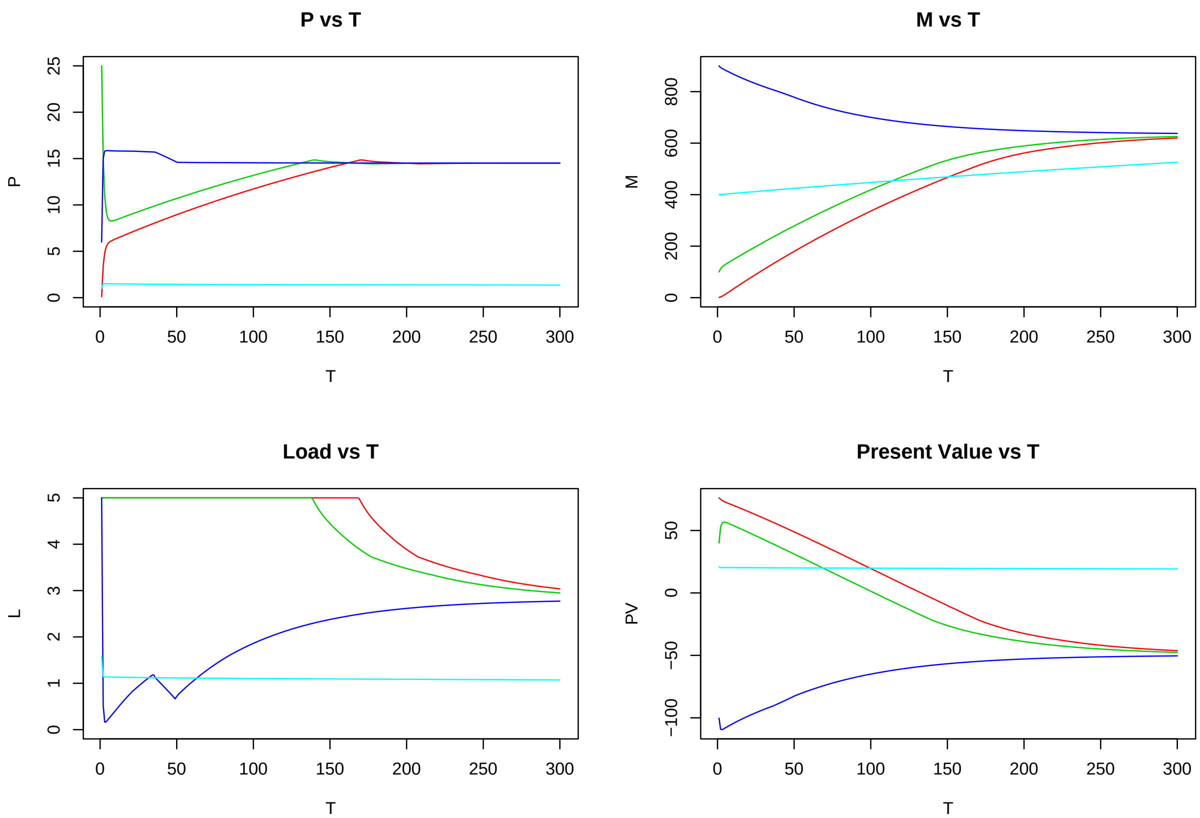

Figure 15: This figure is analogous to Figure 13, but it corresponds to $\delta = .01$.

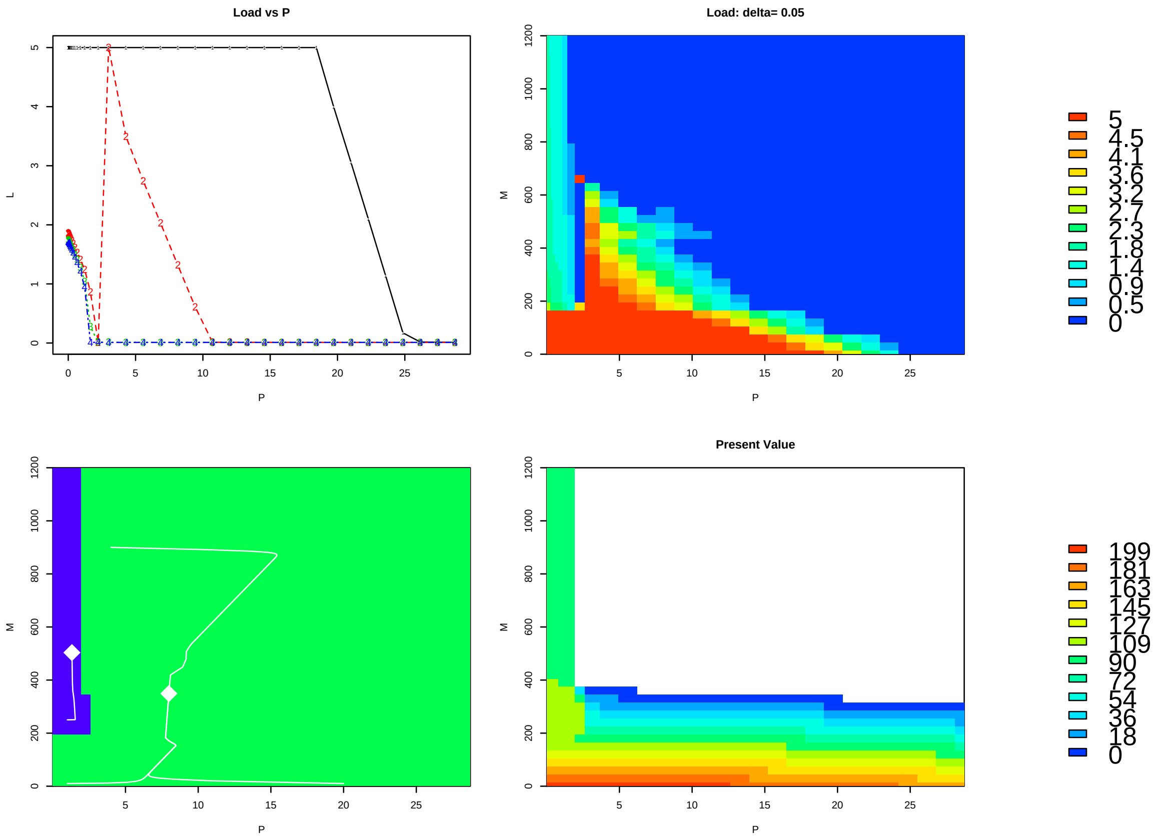

Figure 16: This figure is analogous to Figure 14, but the trajectories correspond to those in the lower left panel of Figure 15. The blue trajectory starts in the upper part of the green area in Figure 15, and the red trajectory starts in the lower left part of the green area. The green trajectory starts in the right part of the green area in Figure 15, and the cyan trajectory starts in the blue region. The time interval used here is too short to show that the cyan trajectory eventually reaches the diamond in the blue region.

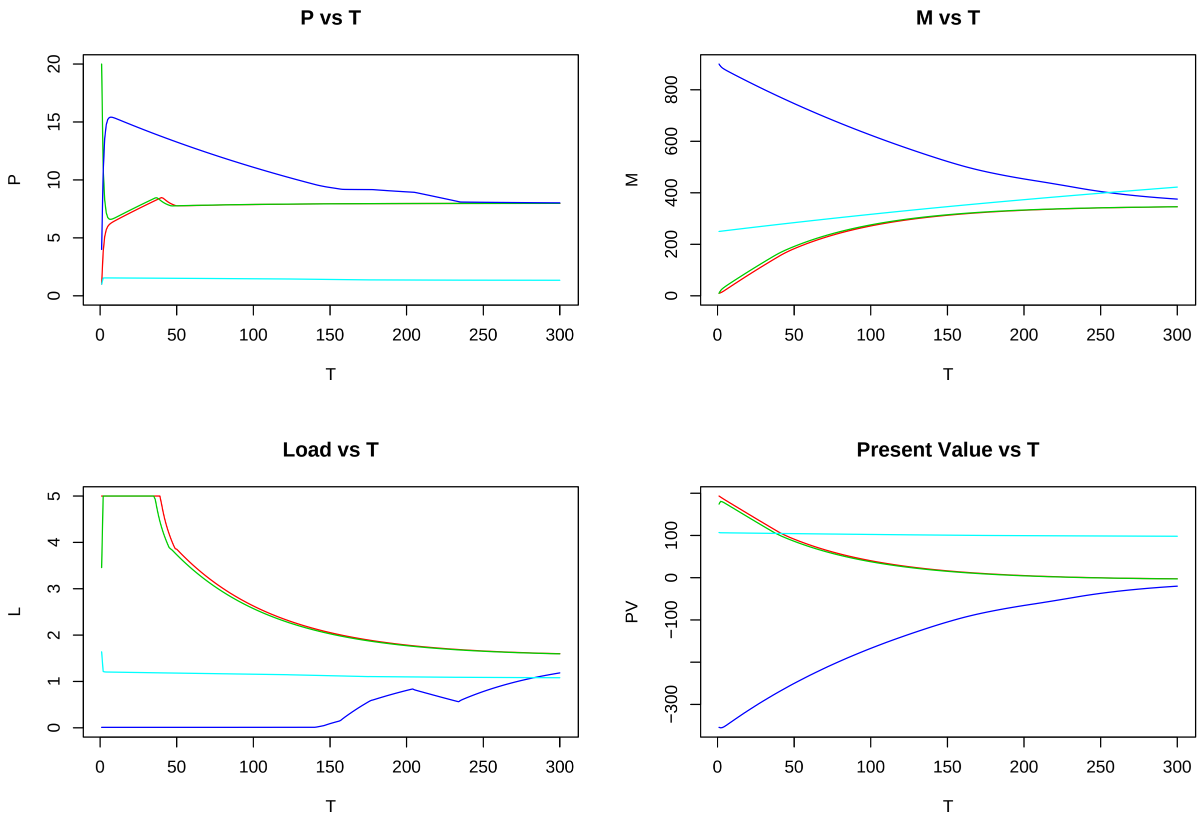

Figure 17: The upper panel shows the mean returns as computed for $M = 240$, $P = 0$ using the two dimensional policy (solid line) and the one dimensional policy (dotted line). The lower panel displays the difference in these mean returns. The scale on the horizontal axis is logarithmic.