

Maximizing return on investment in conservation

William Murdoch^{a,*}, Stephen Polasky^b, Kerrie A. Wilson^c, Hugh P. Possingham^c, Peter Kareiva^d, Rebecca Shaw^e

^aDepartment of Ecology, Evolution and Marine Biology, University of California, Santa Barbara, CA 93108, USA ^bDepartment of Applied Economics, University of Minnesota, St. Paul, MN 55108, USA ^cThe Ecology Centre, School of Integrative Biology, The University of Queensland, St. Lucia 4072, Australia ^dThe Nature Conservancy, 1917 First Avenue, Seattle, WA 98101, USA ^eThe Nature Conservancy, 201 Mission Street, 4th Floor, San Francisco, CA 94105, USA

ARTICLE INFO

Article history: Received 26 March 2007 Received in revised form 22 June 2007 Accepted 6 July 2007 Available online 29 August 2007

Keywords: Conservation planning Cost Benefit Maximizing gain Minimizing loss Species-investment curve Species-area curve

ABSTRACT

Global conservation needs far exceed the available resources, so scarce resources must be used cost-effectively. Although many conservation priory-setting frameworks used by NGO's or public agencies explicitly claim to emphasize efficiency or wise investment, none actually incorporates costs in a formal return-on-investment (ROI) framework. We illustrate here how an ROI framework can be applied to real world resource allocation decisions faced by conservation organizations. We present two examples: (1) allocating resources to purchase land in 21 ecoregions that make up the Temperate Forest Habitat in the US; (2) allocating resources among a variety of conservation actions (not just land purchase) in Mediterranean habitats, with rates of habitat loss factored into the analysis. An important feature of both case studies is that costs vary by orders of magnitude, depending on where or how one is doing conservation. Second, because costs and biodiversity are not well correlated, enormous savings are possible by applying an ROI analysis. Moreover, recommended priorities after including costs in the calculations often deviate substantially from priorities based solely on biodiversity measures. Hence we argue that a major effort of conservationist biologists should be to include and record the costs of conservation actions. If serious attention is not given to returns on investment, it implies that "money is no object."

 $\ensuremath{\textcircled{}^{\circ}}$ 2007 Elsevier Ltd. All rights reserved.

1. Introduction

There is never enough time and money to do everything worthwhile. This is certainly true for conservation where the cost of undertaking important and useful projects far exceeds the available budget. When faced with an array of potential projects and limited means to fund them, standard advice from business and economics is to invest in projects where the rates of return on investment are the highest. Our aims are to show how such a return-on-investment (ROI) approach can be applied to conservation and to illustrate the approach with realistic, if somewhat simplified, examples. Applying ROI to conservation planning should improve the decision-making process, generate better advice, and produce more-effective conservation. It may also be simpler to understand and easier to communicate to donors, decision-makers and the general public than are current approaches.

The academic literature includes hundreds of papers that describe frameworks for priority-setting. Some of these

^{*} Corresponding author.

E-mail address: murdoch@lifesci.ucsb.edu (W. Murdoch).

^{0006-3207/\$ -} see front matter @ 2007 Elsevier Ltd. All rights reserved. doi:10.1016/j.biocon.2007.07.011

frameworks are used by conservation NGO's to guide resource allocation, often by listing the top 25 or top 200 places to work (http://www.biodiversityhotspots.org/xp/Hotspots/;http://www. worldwildlife.org/science/ecoregions/g200.cfm; http://www.bird life.org/action/science/sites/index.html.) Although approaches used by NGOs and agencies seek some form of "efficiency" (e.g., by focusing on areas with high concentrations of endemic species), none explicitly includes costs (Brooks et al., 2006; Groves et al., 2002), and efficiency is typically loosely defined. Without estimates of cost, however, claims of "wise investment" or "efficient allocation of effort" are without support (Naidoo et al., 2006).

Applying return-on-investment analyses to conservation

Return on investment is an intuitively compelling idea with an enormous variety of applications. With a clearly defined and quantifiable objective and limited time, energy, or resources, it makes sense to choose options that maximize the return (defined in units of the objective) per unit investment (dollars, calories, people, time, etc.). In conservation applications, ROI measures the increase in the conservation objective per unit cost of the conservation action. Conservation actions might include land purchase or easements, management of invasive species, fire management, pollution control, lobbying activities, conservation financing aimed at sustainable forestry, and so on. ROI is already used by the many conservation planners who use software, such as MARXAN and SITES, to inform the design of networks of reserves to maximize the number of species protected for a given cost, or to minimize the cost of conserving a fixed percentage of species and/or habitats (Possingham et al., 2000; Sarkar et al., 2006). In ecology, an ROI framework underlies all of optimal foraging theory, which identifies behavioral choices on the basis of maximizing calories gained per unit time or energy expended (Schoener, 1971).

There are several critical steps in applying ROI to resource allocation and planning in conservation (Possingham et al., 2001; Mace et al., 2006).

2.1. Identify a well-defined objective

At the outset we need a clearly stated conservation objective, defined so that it can be measured quantitatively. For example, the objective could be to conserve the maximum number of species, or the value of ecosystem services. (Multiple objectives, such as giving value to both species conservation and the value of ecosystem services, can be accommodated by defining weights for each single objective and maximizing the weighted sum, or by defining a multi-attribute utility function (Keeney and Raiffa, 1976)). In the applications below, we will use number of species conserved as the objective. It is important to note, however, that ROI analysis can be undertaken with *any clearly specified conservation objective* and is not restricted to maximizing the number of conserved species.

Specifying the objective is not a scientific choice. The objective reflects the values, mission, or legal mandate of

the organization producing the conservation plan. It is our experience that organizations often are reluctant to specify a clear objective, perhaps because by leaving objectives vague, disagreements about values or goals do not need to be addressed or resolved. In general, formulation of an objective is the most contentious step in applying an ROI framework to conservation.

2.2. Incorporate realistic estimates of benefits

We next need to estimate how different conservation actions contribute to attaining the objective. For example, how will the number of protected species increase with the addition of a new nature reserve, or with changes in fire management or other management actions?

It is often difficult to assess the benefits of conservation actions because effective monitoring and good data are lacking. In its summary of the current state of affairs on biodiversity conservation, The Millennium Ecosystem Assessment (MEA) stated that "Few well-designed empirical analyses assess even the most common biodiversity conservation measures" (MEA, 2005). At present, estimates of conservation benefits are often just educated guesses or extrapolations based on species-area relationships, though recent work has tried to predict how species conservation changes as a function of how habitats are distributed across the landscape (Cabeza and Moilanen, 2003; Polasky et al., 2005; Moilanen et al., 2005; Nicholson et al., 2006). More work is needed on understanding the biological benefits of conservation actions. All conservation actions are of course based on assumptions about conservation benefits, and inaccurate estimates are not a valid argument against using ROI in particular. Indeed, an ROI framework highlights the importance of conservation monitoring programs and data collection.

2.3. Incorporate realistic estimates of costs

Considering only biological benefits while ignoring costs, or considering costs as a filter only after ranking areas in terms of benefits, results in inefficient use of scarce resources (Ando et al., 1998; Balmford et al., 2000; Naidoo et al., 2006). In current approaches used by NGOs it is common to use biology alone to identify priorities (Brooks et al., 2006). This approach will yield lower returns on investment; for example, sites or strategies with moderate biodiversity levels, but yielding a high return on investment because of very low costs, will be overlooked.

Land purchase costs often vary by orders of magnitude across different potential conservation sites (e.g., Ando et al., 1998; Polasky et al., 2001) so ROI is often more dependent on differences in costs than in benefits (Ferraro, 2003). The gains from including costs can be striking. In a study of African conservation, Moore et al. (2004) found as much as a 66% gain in species coverage when costs were included versus when they were not. At the global level, Balmford et al. (2000) found that up to twice as many species could be conserved for the same budget when costs were included in the analysis (see Naidoo et al., 2006, for a recent summary of this literature). Conservationists frequently believe that estimating costs is the major hurdle to cost-effective conservation, but in fact it is usually easier to estimate costs than benefits (e.g., Naidoo et al., 2006; Wilson et al., 2007). Cost estimates for actions other than land purchase can be generated once the action is defined in sufficient detail.

2.4. Use optimal resource allocation rules rather than ranking schemes in setting priorities and informing decisions

Simple ranking schemes yield fewer benefits per dollar than an ROI approach (Wilson et al., 2006). First, in virtually all cases of interest, conservation actions are not independent of each other, and conservation planning is therefore a "portfolio allocation" problem rather than a simple ranking problem: the benefits or costs of an action depend upon what other actions are taken. For example, the number of species protected by purchase of a particular parcel of habitat may well depend on whether the parcel is connected to other parcels of habitat or is relatively isolated. When returns are not independent across actions, the complete allocation across all actions needs to be considered jointly rather than considering the return on each action in isolation.

Second, ranking schemes give no indication of how much should be invested in the lower- versus higher-ranking areas. It could be that the differences are negligible and essentially everything on the list should receive equal funding. Alternatively it could be that the highest ranking area should receive 70% of the resources and everything else only a small percentage. An ROI analysis provides guidance on differential rates of investment, whereas ranking schemes are really no more than a list of "what is in and what is out."

2.5. Incorporate dynamic considerations into sequential decision-making rather than using static maps for one-time decisions

Consideration of points (1)–(4) is sufficient for a one-time (static) conservation planning decision. Conservation planning, however, is an on-going process in which current decisions set the stage for those to be made in the future (Meir et al., 2004; Costello and Polasky, 2004). Changes in land use, politics and economic conditions alter constraints and opportunities (Armsworth et al., 2006). Climate change and invasions alter the biological landscape. Plans need to change to reflect new realities. Static plans offering the best solution under current conditions are less useful than methods that provide advice for sequential decision-making. ROI analysis can be tailored to fit into either a static or dynamic decision-making framework.

In what follows, we develop two examples to show how ROI analysis may be applied to conservation planning. Our examples are realistic in using real data from areas of conservation interest. The answers, however, are not a set of recommendations. We sometimes simplify for purposes of illustration; in addition, ROI should be used in collaboration with the local conservation actors who know the current facts on the ground.

3. Example 1: Application of ROI to land acquisition for temperate broadleaf and mixed forests in North America

Temperate broadleaf and mixed forests are the dominant habitat type in 21 ecoregions in North America. These ecoregions represent over 2.8 million square kilometers in which 12.8% of the land area has been protected for biodiversity conservation and 31.3% has already been converted to agriculture or development. Much of the remaining forested lands are at risk due to development pressures and habitat fragmentation. Large commercial timber holdings are being sold off at unprecedented rates with a high risk that large forest tracts will be subdivided for developments (Ginn, 2005). Conservation organizations often seek to buy or purchase easements on these lands to restrict subdivision and development or set up sustainable forestry (Rissman et al., 2007). Here we consider the problem of resource allocation among the 21 different ecoregions.

This application is deliberately kept simple to focus on a few key points about cost and allocation of a conservation budget. As a consequence, we exclude many factors important in real conservation decisions. For example, we allow conservation resources to be fully fungible across all ecoregions (i.e. there are no constraints on where the budget is spent). We discuss how to address this constraint and other important issues in the section "Incorporating other factors in ROI analyses" that follows the examples. The primary data in each ecoregion are species richness values (World Wildlife Fund, 2005; Kier et al., 2005) and land values, by county, from the US Department of Agriculture (Table 1). Using the countybased data, we calculated an average land value for each ecoregion, thus ignoring spatial heterogeneity in land values and biodiversity within ecoregions.

Notice that vertebrate and plant richness are only moderately correlated ($r^2 = 0.39$). This illustrates that it is almost never possible to maximize conservation of all biodiversity components simultaneously.

3.1. The simplest ROI analysis

In all analyses in this paper we tally the number of species protected in each ecoregion, without regard to whether some or all of these species are already protected in a different ecoregion (i.e. we ignore complementarity). This simplification is commonly made when species lists by region are absent, which is true for many regions of the world for many taxonomic groups, or where resources are allocated over large scales where species overlap is relatively small. We do have species lists for each US ecoregion and could take into account species identity and therefore the complementarity of different ecoregions. We discuss the more complicated problem of species identity and complementarity at the end of the paper.

As the simplest possible illustration of ROI, suppose that conserving the same area of land in each ecoregion protects the same fraction of species in each ecoregion; for example, assume that conserving 100,000 acres in any ecoregion conserves 50% of the species in that ecoregion (Table 3 lists areas in ecoregions). In this case, we calculate ROI by:

Ecoregion	Land value (\$/acre)	Vertebrate species richness	Vertebrate richness ranking	Plant species richness	Plant richness ranking
Piedmont	1915	483	1	3363	1
Upper East Gulf Coastal Plain	830	453	3	3363	1
Cumberlands and Southern Ridge Valley	1215	444	4	2487	3
Western Allegheny Plateau	1248	401	14	2487	3
Central Appalachian Forest	1624	415	10	2398	5
Southern Blue Ridge	2234	404	12	2398	5
Interior Low Plateau	1184	440	6	2332	7
Ozarks	870	436	7	2332	7
North Central Tillplain	1615	388	16	2243	9
High Allegheny Plateau	1281	360	19	1883	10
Ouachita Mountains	822	419	9	1743	11
Lower New England/Northern Piedmont	5606	420	8	1695	12
North Atlantic Coast	9644	405	11	1695	12
Northern Appalachian-Boreal Forest	1100	366	18	1496	14
Chesapeake Bay Lowlands	2543	442	5	1488	15
Mississippi River Alluvial Plain	1087	464	2	1468	16
Great Lakes	1464	403	13	1459	17
Superior Mixed Forest	867	377	17	1459	17
Prairie-Forest Border	1414	389	15	1420	19
St. Lawrence/Champlain Valley	918	351	20	1381	20
Willamette Valley	3126	324	21	1067	21

Table 1 – Species richness and land value for 21 Temperate Forest ecoregions in the US. Ecoregions are ranked by plant and vertebrate species-richness

The species richness data were drawn from WWF ecoregions, which in some cases have slightly different boundaries than TNC ecoregions. The numbers in this table were extrapolated to TNC ecoregions by producing a GIS overlay of WWF and TNC ecoregions and calculating an area-weighted average.

- calculating the benefit, which is the number of species conserved by buying 100,000 acres (assumed to be 50% of the ecoregion's species);
- calculating the cost, which is 100,000 acres times the land value per acre in the ecoregion;
- dividing the benefit by the cost.

Applying these calculations for plant and for vertebrate species generates an ROI for the two taxonomic groups by ecoregion (Table 2). Note that in this simple example, we have analyzed a single level of investment rather than how ROI changes with levels of investment in any ecoregion.

Even this very simple analysis immediately demonstrates several useful points. First, priority rankings are different when using only species richness (benefits alone) versus using ROI (benefits and costs). Piedmont is the highest ranked ecoregion in terms of species richness for both vertebrates and plants, but only 15th in terms of vertebrates conserved per \$1 million and 7th in terms of plants conserved per \$1 million. Superior Mixed Forest is 17th in terms of vertebrate richness but 4th in terms of vertebrates conserved per \$1 million. This illustrates the danger of first ranking by biodiversity and then applying a cost filter: in a list of 21 ecoregions ranked by biodiversity, Superior Mixed Forest would have been culled from any plausible short list of ecoregions to work in. On the other hand, no matter how you slice it, the Upper East Gulf Coast Plain is a high priority, ranking 1st in terms of both vertebrates and plants conserved per \$1 million, tied for 1st in terms of plant richness and 3rd in terms of vertebrate richness.

Second, the example shows the fundamental role played by costs. Costs vary by more than an order of magnitude, from a low of \$822 per acre in the Ouachita Mountains to a high of \$9644 on the North Atlantic Plain. By contrast, vertebrate species richness varies by less than 50% and plants by threefold. The greater variation in costs makes it of great importance to include costs from the beginning. The top three ecoregions ranked in terms of ROI for plants or vertebrates are those with the lowest land values. Two of the three (Ozarks and Ouachita Mountains) have only intermediate rank in terms of biodiversity.

3.2. Adding diminishing returns to the ROI analysis

The preceding analysis treated the return on investment in each ecoregion as a constant when, in reality, we expect a lower marginal conservation return for each *additional* investment in an ecoregion as the baseline area protected increases. In particular, we accumulate many species when we first begin to conserve habitat, and with further conservation of habitat we accumulate fewer additional novel species. Here we undertake an incremental ROI analysis that takes this fact into account using the concept of a species–area curve, and apply the concept to plant diversity. This approach incorporates information about the amount of land already protected within an ecoregion.

The relationship between species richness, S, and area, A, can be described by the familiar equation for a species–area curve, $S = \alpha A^z$, where α is a constant and z < 1, so that species are accumulated more slowly as area increases. An estimate of z is crucial if one wants to predict the impact of habitat protection on diversity. Values of z are typically estimated to be between 0.15 and 0.4 depending on the type of habitat, as well as the distribution and dispersal abilities of the species under consideration. Larger z values are typical of true oceanic is-

Table 2 – ROI by ecoregion				
Ecoregion	Vertebrates conserved per \$1 million	Ranking	Plants conserved per \$1 million	Ranking
Upper East Gulf Coastal Plain	2.73	1	20.26	1
Ozarks	2.51	3	13.40	2
Ouachita Mountains	2.55	2	10.60	3
Cumberlands and Southern Ridge Valley	1.83	8	10.23	4
Western Allegheny Plateau	1.61	10	9.96	5
Interior Low Plateau	1.86	7	9.85	6
Piedmont	1.26	15	8.78	7
Superior Mixed Forest	2.17	4	8.41	8
St. Lawrence/Champlain Valley	1.91	6	7.52	9
Central Appalachian Forest	1.28	14	7.38	10
High Allegheny Plateau	1.41	11	7.35	11
North Central Tillplain	1.20	16	6.94	12
Northern Appalachian-Boreal Forest	1.66	9	6.80	13
Mississippi River Alluvial Plain	2.13	5	6.75	14
Southern Blue Ridge	0.90	17	5.37	15
Prairie-Forest Border	1.38	13	5.02	16
Great Lakes	1.38	12	4.98	17
Chesapeake Bay Lowlands	0.87	18	2.93	18
Willamette Valley	0.52	19	1.71	19
Lower New England/Northern Piedmont	0.37	20	1.51	20
North Atlantic Coast	0.21	21	0.88	21

lands and for species with relatively limited movement, whereas smaller values are common for highly mobile species and for pseudo-islands that are actually patches of habitat embedded within a continuous landscape (such as nature reserves). For the example worked here, we use a z = 0.2, which has commonly been reported for plants in mainland habitats.

Fig. 1 shows species–area curves for five of the US ecoregions. Each curve starts out steeply, indicating a rapid initial increase of protected species per unit area of acquired land, and then one obtains progressively smaller gains as additional lands are protected and the curve gradually flattens out. If we know the area currently protected in the ecoregion, we can then determine the incremental benefit with each addition to the total protected area. For the equation used in Fig. 1, the marginal increase in species protected is the derivative of the species–area equation with respect to A, $\frac{dS}{dA} = z\alpha A^{z-1}$. Thus, the marginal gain in species declines as A increases. Note that such estimates do not address the question of population viability: the number of species protected in an isolated small area is likely to be overestimated because some "expected" species populations will not be viable.

If we have an estimated species–area curve, and a cost per unit area of land, we can easily produce a *species–investment curve*, where the area conserved has been translated into a dollar value (Fig. 2) (Naidoo and Adamowicz, 2005). The species–investment curves indicate how many species could be conserved if land has to be purchased at the average land value in the ecoregion. Ecoregions with expensive land are adjusted downward relative to ecoregions with inexpensive land (e.g., North Atlantic Coast falls below Superior Forest in Fig. 2 while being well above it in Fig. 1).

Using graphs such as those illustrated in Fig. 2 for every ecoregion in our analysis, we can calculate the increase in



Fig. 1 – Species–area curves for five US ecoregions. The number of plant species in the ecoregion, S, increases with area, A, following the standard formulation $S = \alpha A^z$ with z = 0.2; we solved for α using the plant data from each ecoregion.





species conserved per million dollars invested. Table 3 shows this measure for each ecoregion, evaluated by starting from the current amount of protected area within each ecoregion. Clearly, the starting point (the amount of area already protected) influences the return on investment one obtains.

The ecoregions that rank highly in this analysis have relatively little existing protected area (which makes incremental returns high) and tend also to be highly ranked in the analysis shown in Table 2 (i.e., have high species-per-dollar rankings). The Upper East Gulf Coast continues to rank highly, as does the Ouachita Mountains. The North Central Tillplain jumps to the top of the list because none of it is currently protected. Protecting some area in the North Central Tillplain is extremely valuable in this context, precisely because no area is currently protected and the marginal increase in species conserved with an increase in protected area is therefore high. In reality, this benefit is no doubt overestimated because there is some protected area, in local parks for example, and, as noted, the species-area curve may overestimate population viability in small areas. Other ecoregions move up or down the list relative to Table 2 depending upon how much area is currently protected within the ecoregion. The marginal ROI for each ecoregion will decline as investment proceeds. For example, as shown in Fig. 2, the first \$10 million invested

Table 3 – Marginal ROI analysis						
Ecoregion	Plant richness	Total area (million hectares)	Current protected area (ha)	Exponent, α	Land value (\$/ha)	Increase in plant species conserved per \$1 million investment
North Central Tillplain	2243	12.316	0	85.65	3990.75	258.54
Upper East Gulf Coastal Plain	3363	13.708	49,348	125.70	2050.97	2.15
Ouachita Mountains	1743	4.647	29,742	80.88	2031.21	2.09
Western Allegheny Plateau	2487	10.782	33,425	97.53	3083.88	1.51
Piedmont	3363	17.137	25,705	120.21	4732.07	1.50
Cumberlands and Southern	2487	12.566	35,185	94.59	3002.33	1.45
Ridge Valley						
High Allegheny Plateau	1883	6.826	34,131	80.91	3165.42	1.20
St. Lawrence/Champlain Valley	1381	6.146	42,409	60.60	2268.43	1.06
Willamette Valley	1067	1.485	10,097	62.21	7724.51	1.00
Chesapeake Bay Lowlands	1488	4.396	22,418	69.82	6283.89	0.73
Interior Low Plateau	2332	19.329	110,175	81.37	2925.73	0.51
Ozarks	2332	13.892	219,500	86.93	2149.82	0.43
Mississippi River Alluvial Plain	1468	10.972	141,538	57.37	2686.04	0.32
Central Appalachian Forest	2398	9.656	207,604	96.14	4012.99	0.27
Prairie-Forest Border	1420	15.833	134,580	51.57	3494.07	0.23
Southern Blue Ridge	2398	3.809	276,547	115.79	5520.33	0.19
Lower New England/Northern	1695	9.401	50,763	68.32	13852.73	0.17
Piedmont						
North Atlantic Coast	1695	5.138	39,561	77.09	23830.84	0.14
Superior Mixed Forest	1459	20.760	1,123,097	50.19	2142.40	0.07
Northern Appalachian-Boreal Forest	1496	33.454	1,361,569	46.78	2718.16	0.04
Great Lakes	1459	57.519	1,000,831	40.94	3617.62	0.04

in North Central Tillplain protects almost 410 species, whereas the next \$10 million protects an additional 60 species. Allocating a conservation budget across ecoregions, taking into account changing incremental ROI values, requires a slightly more sophisticated analysis, which we cover next.

3.3. Optimal allocation of a conservation budget across ecoregions

We extend the approach of the previous section to provide the result we have been building towards: the optimal allocation of a fixed budget across all ecoregions, where the objective is to maximize the total number of species conserved. We assume, again, that the number of species conserved by ecoregion, summed over all ecoregions, equals the number of species conserved overall. A more sophisticated analysis would take account of which species are conserved in each ecoregion and place higher value on conserving species not conserved elsewhere (see below).

Allocation of resources depends on the total funds available, and we show this by considering various budget levels, starting at \$100 million and going up to \$500 million in \$100 million increments (Table 4). At the initial budget level, investments are made in only 6 of the 21 ecoregions, with large investments made in only three: North Central Tillplain, Upper East Gulf Coastal Plain, and Ouchita Mountains. The North Central Tillplain gets the most investment in part because it begins with no protected area, making initial investments in conservation highly productive. Both the Upper East Gulf Coastal Plain and the Ouchita Mountains rank highly because they have relatively little protected area, high species richness, and low land values. As the budget is increased, investments are made in more ecoregions (9 out of 21) and now the ecoregion getting the most investment is the Upper East Gulf Coastal Plain. If the budget were expanded further, more ecoregions would receive investments, and there would be continued expansion in those ecoregions already receiving investments. The total amount of money available can dramatically shift the proportional resource allocation patterns among ecoregions (Table 4). Thus, when there is only \$100 million to spend, North Central Tillplain receives 37.7% of the budget; but when there is \$500 million to spend, North Central Tillplain receives only 16%.

Clearly, conservation is not as simple as an application of species-area curves, land costs, and an optimization algorithm. To succeed in the real world, conservationists have to deal with public attitudes, possibility of government funding, landowners unwilling to sell, the need to raise money around some compelling project, and threats that cannot be addressed simply by land acquisition (e.g., invasive species). Nevertheless, for all its simplifications, this example illustrates the importance of evaluating information that is rarely explicitly examined by conservation planners, and doing so in an analytical framework. Conservation planners routinely quantify benefits (some measure of species or habitats protected), and routinely apply species-area curves. But economic costs rarely, if ever, find their way into prioritization schemes, and the effect of initial conditions (amount already protected) is not analyzed quantitatively. For example, a recent review of nine different global prioritization frameworks for allocation of funds revealed that none explicitly includes costs and that, although initial conditions may be considered, that consideration does not entail a formal application of species-area curves and diminishing returns as the area protected increases (Brooks et al., 2006).

4. Example 2: Application of ROI to a range of conservation actions in the presence of threats

This example explores the allocation of funds among 17 Mediterranean ecoregions in Australia, Chile, California, and South Africa (Table 5) (see Wilson et al., 2007). Mediterranean habitats are rich in species, highly degraded throughout much of the world, and a priority for many conservation organizations (Cowling et al., 1996; Myers et al., 2000; Hopper and Gioia, 2004). Our question then is, how should money be allocated among different places and different activities, given the on-going loss of habitat?

We aim to solve here the problem, familiar to conservationists, of trying to choose between saving the most diverse habitats and the most threatened. If we ignore threat intensity, then in Year 1 we may protect a low-risk habitat that yields the highest number of species protected per dollar spent; in the meantime, however, we may have lost the opportunity to protect a high-risk habitat where some irreversible change has occurred. By contrast, if we had first saved the high-risk habitat in Year 1, the low-risk habitat would likely still be available for action in Year 2 or even later. (For those who like analogies: if we want to save all lives at an accident scene we should work first on those most likely to die.) Sites at higher risk, of course, may also be more costly.

Example 2 expands on Example 1 in two fundamental ways.

 Multiple threats and multiple actions: Land purchase is only one of a number of actions we can take to protect biodiversity, and in many cases it is neither feasible nor appropriate. Here we begin by considering the threats to biodiversity in each ecoregion and then consider a range of conservation actions to abate the threats.

Table 5 – Ecoregions ranked by plant species richness					
Ecoregion	Plant and vertebrate species richness				
Montane fynbos and renosterveld	6805				
Lowland fynbos and renosterveld	3468				
Esperance mallee	3281				
Chilean matorral	2806				
Swan Coastal Plain Scrub and Woodlands	2658				
Jarrah-Karri forest and shrublands	2487				
Interior chaparral and woodlands	2473				
Montane chaparral and woodlands	2421				
Eyre and York mallee	2383				
Mount Lofty woodlands	2393				
Southwest Australia woodlands	2338				
Southwest Australia savanna	2357				
Coolgardie woodlands	2143				
California coastal sage shrub	2088				
Naracoorte woodlands	1874				
Albany thickets	1616				
Murray-Darling woodlands and mallee	1600				
The five according receiving most funding in the POI englysic are					

The five ecoregions receiving most funding in the ROI analysis are in bold.

2. Dynamics of landscape and threats: Example 1 assumes that all habitats will be available for action whenever we decide to act, and will be lost if action is not taken. In reality, however, habitats or species are continually being lost as a result of human activities, but not all face certain destruction within a given period. Our allocation of resources takes account of these anticipated future losses. We incorporate into the analysis a predicted loss of unprotected habitat in each time period, and we derive the expected number of species lost using the species–area relationship.

The objective (as in Example 1) is to maximize the expected number of species remaining at the end of the planning period or, stated otherwise, to minimize the expected loss of species. Taking explicit account of the threat of habitat loss places a premium on ecoregions that face high threat of loss, in addition to consideration of biological value and cost. A static analysis that maximizes the number of species in the set of protected sites can generate results that are inferior to approaches that minimize the expected number of species lost, taking explicit account of threats (Costello and Polasky, 2004; Meir et al., 2004; Wilson et al., 2006). The difference between approaches arises because not all habitats outside of protected areas will be lost, as is implicitly assumed in the static approach (Costello and Polasky, 2004).

Ecoregions differ in both the type and the intensity of threats. We use a generic measure of human activity, the "Human Footprint," developed by the Wildlife Conservation Society (Sanderson et al., 2002), to predict the intensity of each threat in each ecoregion, and hence the area of habitat that will be lost owing to each threat if no action is taken (see Supplementary material).

4.1. Implementation

Let's assume we have \$100 million to spend each year over 20 years. The objective is to allocate these funds to minimize the expected loss of species (i.e., maximize the number of species conserved) across all ecoregions at the end of 20 years. As in Example 1, we allocate funds to the ecoregion(s) with the highest benefit to cost ratio, but now the benefit is the expected number of species that we can prevent from being lost.

We use a simple rule, termed "minimize loss" (min loss), to achieve this objective (Wilson et al., 2006). Each year, we calculate the marginal benefit from investing in a particular action in a particular ecoregion. The marginal benefit is the number of species that can be saved per dollar, species that would otherwise be lost if no action were taken. We then allocate funds to the set of conservation actions that minimizes the expected loss per dollar, until the annual budget has been exhausted. Because this simple rule does not consider the full set of options through time but looks only at a single period, it does not necessarily result in an optimal solution. However, the min loss rule has been shown to find optimal or near-optimal solutions in small problems (Costello and Polasky, 2004) and to perform well when there are substantial differences in threat intensity among ecoregions (Wilson et al., 2006).

To apply the min loss rule in an ROI analysis we need to know, for each ecoregion: the threats, the mitigating actions available to abate each threat, their cost per unit area, the area over which each conservation action is required and the area already receiving each action; the biodiversity benefit obtained by abating each threat; and the predicted background rate of species loss. Where the cost is annual and recurring, (for example, feral predator control) we calculate the cost to endow the action permanently. For this example, the threats, mitigating actions, and associated cost for each threat that we evaluate are in Table 6 (see Wilson et al., 2007 for details on how these were obtained).

We estimated the potential biodiversity benefit from each mitigating action in each ecoregion by using the IUCN red list to determine the proportion of vertebrate and plant species for which the threat being considered was listed as a major source of extinction risk. This proportion was then multiplied by the total number of vertebrate and plant species in the ecoregion as identified in the WWF Wildfinder database and by Kier et al. (2005) respectively. Since there are more vascular plants than vertebrates in Mediterranean ecoregions, the biodiversity benefit is dominated by plant species. The number of at-risk species for each action–ecoregion combination is in Supplementary material.

We have assumed that investing in an action aimed at a specific threat will save the number of species otherwise expected to be lost to that threat. That is, we have not dealt with the fact that some species face multiple threats.

4.2. Optimal allocation of resources across ecoregions and conservation actions

The expected number of at-risk species saved per dollar invested is calculated in Year 1 for each of the 51 action–ecoregion combinations. The algorithm then prioritizes spending that year on actions that deliver the biggest conservation bang-for-buck. These calculations are repeated each year for 20 years. Fig. 3 shows the sequence of resource allocations over the first 5 years.

In Year 1, invasive predator control in Swan Coastal Plain Scrub Woodlands in Australia and the control of priority noxious weeds on public lands in the Californian coastal sage shrub receive the largest allocations of funds (Fig. 3). The rates of habitat loss in these ecoregions are high and the area of remaining habitat small. Because predator control (even though it is endowed over 20 years) is cheap (about \$7000/km²), the marginal return on investment is very high: in the Swan Coastal Plain, applying predator control to only 4600 km² protects 143 threatened species at a cost of about \$32 million. Fire management in the Chilean matorral also has a very high ROI: the cost is only about \$500/km², and 500 species are predicted to be protected when this action is applied to all the available area at a total cost of only about \$12 million. The gains from weed control in California coastal sage are not fully realized until another large investment is made in Year 2 (Fig. 3).

Allocations shift over time for several reasons. First, there is a limit to the area available for a particular action in any ecoregion (e.g., almost all the area available for predator control in Swan Coastal Plain Scrub Woodlands is covered in Year 1). Second, the marginal benefit to cost ratio of a particular action declines as investment proceeds (due to diminishing returns), so actions with initially lower benefit to cost ratios typically become relatively more cost-effective as the first actions taken yield diminishing returns. Third, loss of habitat due to human actions further limits the land available for conservation and modifies the marginal benefit to cost ratio.

Fig. 4 (open bars) shows the cumulative amount of money allocated to various actions in order to minimize the expected loss of species across all ecoregions after 20 years. Of the 51 possible actions in the 17 ecoregions, only 16 receive funds over 20 years. At the end of 20 years, the greatest amount of money is allocated to invasive plant control in the Chilean matorral (even though this action is not funded in the first 5 years). The rate of habitat loss in this ecoregion is high, the benefit to cost ratio of this action is high, and currently this conservation action is being applied over a limited area. We discuss the solid bars in Fig. 4 below.

Table 5 shows that the five ecoregions (bolded) that receive the most investment after 20 years have intermediate species richness; the three ecoregions with the greatest species

Mediterranean region	Threat	Action	Cost (US\$ per km ²)
Australia	Habitat fragmentation Introduced predators (cats and foxes) Phytophthora cinnamomi	Revegetation Invasive predator control and research Phytophthora cinnamomi management	300,000 7000 (endowed) 515,000 (endowed)
Chile	Invasive plants Altered fire regimes Conversion of natural habitat	Invasive plant removal, herbicide application and revegetation Fire suppression Land acquisition and management	126,757 516 (endowed) 277,273
South Africa	Conservation of natural habitat Invasive plants	Land acquisition and management Invasive plant removal, herbicide application, and follow-up treatment	53,000 92,900
California	Invasive plants Invasive plants Conversion of natural habitat Altered fire regimes Altered fire regimes	Control of priority noxious weeds on public land Control of riparian invasives Land protection Fire suppression Fuel reduction	3,300,000 4,447,897 1,013,882 1,633,358 526,765

Table 6 – Threats, mitigating actions, and associated costs for each threat in four Mediterranean ecoregions



Fig. 3 - Investments in each mitigating action (US\$) during the first five years.



Fig. 4 – Total investments over 20 years. Open bars are the results of the minimize-loss approach pursued in this paper in Mediterranean ecoregions; solid bars are results obtained for the same data using the maximize-gain approach (see Supplementary material for a list of the ecoactions designated by the codes on the X-axis).

richness receive no allocations over 20 years. In conventional resource allocation frameworks, these three ecoregions with the greatest species richness would invariably be selected as high priority regions for conservation investment. That conclusion is clearly unwarranted when one adopts the ROI approach. The investment schedule is determined by the area requiring investment, the benefit to cost ratio and the urgency of the investment.

The main messages of Example 1 are reinforced: the most species-rich ecoregions do not necessarily yield the best return on investment, and ranking indices underperform relative to ROI analysis.

4.3. Comparison with maximize-gain approach

Wilson et al. (2007) analyzed exactly the same data as we did above, but ignored ongoing loss of habitat and instead applied the maximize-gain approach of Example 1; their allocations are represented by the solid bars in Fig. 4. While the two approaches generate somewhat similar allocations, there are some striking differences.

The maximize-gain approach allocates the most to land acquisition to abate agricultural conversion in the Montane and Lowland Fynbos ecoregions of South Africa, because the ecoregion contains large areas of arable land that are unconverted but largely unprotected, and the biodiversity benefit is high and the cost comparatively low (Wilson et al., 2007). But this ecoaction receives no allocation in 20 years under the minimize-loss because the predicted rate of habitat loss in this ecoregion is zero. That is because the human population in South Africa is predicted to decline, and hence the Human Footprint calculation predicts no loss of habitat. (In reality, some Fynbos habitat will no doubt be converted to agriculture or urban use, which our analysis fails to capture; see Section 7).

5. Minimize loss versus maximize gain

Experienced NGO conservationists tend to fall into one of two camps in setting priorities: (1) invest where we can gain the most diversity or conservation value, versus (2) invest where there is the greatest imminent threat, i.e. places we most expect to lose in the near future. Scientists and managers argue about these two options, and develop top 10 or top 20 lists on the basis of both approaches. These two views have formal counterpoints in ROI analysis: maximize gain versus minimize loss. Our Example 1 followed the first approach, Example 2 followed the second. A major value of formal ROI analyses is that the similarities and differences between these two priority-setting approaches become transparent, as we illustrate in this section.

There is no fixed rule mandating when one should pursue a minimize-loss rather than a maximize-gain approach to conservation. In general, the minimize-loss approach is better (i.e. closer to the true optimal solution) when rates of loss of species vary markedly from one place to another (as in Example 2), provided the losses can be both estimated, and stemmed by conservation action. The maximize-gain approach is preferable when loss rates are low and uniform.

We emphasize the differences between the approaches with a simple hypothetical example of two ecoregions in which species losses are caused by habitat destruction, and the only feasible conservation action is land purchase (Fig. 5). As in Fig. 2, each species–investment curve is derived from a species–area curve, which in this case would depict the number of at-risk species accumulated as the amount of habitat exposed to the particular threat increased: the righthand end of that curve would depict the total number of species present in the at-risk habitat. In Fig. 5, the upper ecoregion (Ecoregion 1) is more species-rich. We assume land costs are the same in the two ecoregions, so Ecoregion 1 also has more habitat remaining and its total purchase would require a larger investment (the X-axis), and we assume no existing land conserved in either ecoregion.

Suppose we are at the start of Year 1 and can spend \$100 million. First, the maximize-gain strategy starts at the origin of Fig. 5 and determines, for each ecoregion, how many species will be saved by the first investment of \$100 million. The answers are shown by the vertical line at \$100 million



Fig. 5 – Contrasting maximize-gain (top graph) with minimize-loss (bottom graph) ROI approaches for two hypothetical ecoregions. See text for explanation of the two different approaches, and why they indicate different priorities.

in Fig. 5 (top): we would protect 181 species in Ecoregion 1 (ROI = 1.81 species per \$1 million) and 131 species in Ecoregion 2 (ROI = 1.31 species per \$1 million). This strategy tells us to invest our \$100 million in Ecoregion 1.

For the minimize-loss strategy we need to know the rate of loss caused by habitat destruction. Let us assume the rate is twice as high in Ecoregion 2 so, if we did nothing to prevent it, this process would destroy in the next year \$100 million worth of Ecoregion 2 habitat and \$50 million worth of Ecoregion 1 habitat. If we did nothing to prevent the losses, the right-hand end of each curve in Fig. 5 (bottom) would move left the distance indicated by the horizontal dotted line. Thus, in the minimize-loss strategy, we start at the right-hand end of the curve, rather than the origin. The consequent loss of species in each ecoregion is indicated by the vertical dotted line at the new potential endpoint of the curve. Under the minimize-loss approach the ROI is the number of species whose extinction each dollar can be expected to prevent. In this case we maximize ROI by investing in Ecoregion 2 rather than Ecoregion 1. We can save 34 species by spending \$100 million in Ecoregion 2, giving a benefit:cost ratio of 0.34 species per \$1 million, compared with saving 13 species by spending \$50 million in Ecoregion 1, which has a benefit to cost ratio of 0.26 species per \$1 million (see the dotted horizontal and vertical lines in Fig. 5). Ecoregion 2 is preferred for investment under the minimize-loss strategy for two reasons. First, even though it is less species-rich, Ecoregion 2 is currently in a steeper part of the species-investment curve (because there is less habitat left), so the marginal gain from each dollar invested is greater. Second, Ecoregion 2 faces a higher rate of potential habitat destruction, so there is more urgency to protect the habitat remaining.

In general the two approaches will differ most when rates of habitat loss are high and heterogeneous among ecoregions. In those cases it matters a great deal whether you care most about achieving the maximum gain or about achieving the minimum loss. However, if rates of habitat loss are either low, or relatively homogenous among different ecoregions, then it does not really matter which prioritization model one operates under.

6. Incorporating other factors in ROI analyses

To make central points clearly, the examples we have used in previous sections exclude some aspects of conservation decision-making in the real world. Here we indicate briefly how to expand the analysis to incorporate additional important factors.

Complementarity and endemism: Because different ecoregions may share some species in common, a given species may be protected in multiple ecoregions. If the goal is protect the largest number of different species in at least one ecoregion, we need to make use of the fact that different ecoregions may have complementary sets of species, which requires us to know the identity, not just the number, of species. This problem is familiar in reserve design, and is handled by software such as MARXAN and SITES (Possingham et al., 2000). Optimal solutions providing maximum coverage of species that incorporates complementarity can be found even for relatively large problems (see, for example, Ando et al., 1998; Church et al., 1996; Csuti et al., 1997). Complementarity can be incorporated into the ROI approach by considering the number of species that occur in multiple ecoregions separately from those that are endemic to an ecoregion, and representing this information using separate species-area curves. For example, if we have two ecoregions A and B, we need to consider not only an 'A species-area curve' and a 'B species-area curve' for species endemic to ecoregions A and B respectively, but also an 'AB curve' for the species found in both ecoregion A and ecoregion B. The number of species-area curves thus increases combinatorially as the number of ecoregions increases. The ecoregion that will be favored for investment is the one where for each dollar invested, the sum of the additional species that would be protected across all of the species-area curves is greatest (Underwood et al., unpublished manuscript).

Investment risk: Investment risk is the probability that a "protected" species or habitat will in fact be lost in the future, for example because of inadequate enforcement of laws (e.g., illegal logging) or because a future government decides to renege on prior agreements. One way to include investment risk is to "discount" the conservation benefit. For example, we might expect to conserve 100 plant species by protecting 10,000 ha of forest in Sumatra. But if we think that, over the next 50 years, it is likely that 5000 ha will be illegally logged, we could reduce the expected area conserved to 5000 ha and, using a species area curve, reduce the expected benefit accordingly, to say 70 species conserved instead of 100. Alternatively, if the analysis considers particular sites, we can introduce a loss rate of protected sites that reflects local vulnerability.

Funds not fully fungible: We assumed there was no restriction on where conservation resources could be allocated. In US NGOs, however, funds are typically raised by and spent in each state so that funds are not fully fungible across ecoregions. It is straightforward to incorporate constraints that limit where funds can be spent. For example, in the case of funds raised by state chapters that must be spent within that state, there would be a constraint requiring at least a minimum level of conservation expenditure within each state in addition to the overall budget constraint.

Leverage: Investment may stimulate different amounts of matching funds in different regions. Matching funds effectively change the cost of making certain investments. For example, if funds are matched on a one-for-one basis then the cost to the NGO is reduced by 50%. More complex matching formulas can also be incorporated into the analysis.

Start-up costs: There are costs to starting a new country or regional program. A way to incorporate such costs into the analysis is to add them to the first unit of conservation effort allocated to the country or region. Start-up costs add some complexity to the analysis. Even though there often will be high returns after the initial investment, we may need to do the analysis with and without the new country or program to determine if it makes sense to incur the start-up cost.

Intangible benefits: NGOs often carry out conservation actions in part because it will encourage, empower, or guide potential or actual local partners, or will lead to an improved political climate. The only rational basis for such actions is that they are expected to lead to greater benefits (e.g. more species saved) or reduced costs (or investment risk) in the future. These conservation actions could be included in an ROI framework as a form of investment that builds capacity that will yield future benefits. We still need, however, some estimate of the future benefits.

Unpredictable loss of available sites: Because our examples conserve at the level of the ecoregion, we simply protected given amounts of land whose actual positions in the ecoregion were unspecified. A new problem arises when we consider actual sites within ecoregions: although we can often predict the approximate rate at which sites will be lost to human development, we cannot necessarily predict which sites will be lost. The formal machinery for calculating the optimal budget allocation in this situation is stochastic dynamic programming (SDP), which allocates funds each year while taking into account all possible future states of the system (Costello and Polasky, 2004; Wilson et al., 2006). However, in most real conservation situations, SDP cannot be used because there are too many sites and too many possible patterns of potential loss. Fortunately, simple rules that can be calculated quickly, such as min loss, have been shown to generate optimal or near-optimal results in such problems (Wilson et al., 2006).

7. Discussion

The ROI approach to conservation planning is a decision-support tool that provides practical guidance on how to allocate funds across many potential projects or actions. Equally important, this approach also can and should be used to explore alternative scenarios and for systematic learning.

In most real-world conservation decisions, unlike our somewhat simplified examples, there will be large uncertainties about the benefits and the costs of conservation actions. It is important to realize that uncertainty should not be used as an argument against the application of ROI. Conservationists already apply sophisticated planning tools on the basis of highly uncertain data, such as predicted species occurrences as opposed to observed occurrences (Polasky et al., 2000; Camm et al., 2002; Wilson et al., 2005; Moilanen et al., 2006). Further, because the costs of conservation actions often vary by two or three orders of magnitude across available options, it will typically be better to accept estimation errors of even 50% than to exclude cost considerations.

In general, the most fundamental challenges facing ROI analyses stem, not from estimating costs, but from incorporating threats and averted species losses, and accurate assumptions about the effectiveness of conservation actions. For example, the current ROI framework assumes that if a species is protected by removing one threat, it is not going to be lost due to other threats. This assumption tends to overestimate the number of species saved by a particular investment when there are multiple threats. It is also difficult to predict future rates of habitat loss. The data we have used as a proxy for habitat loss rates are obviously too generic for this purpose; to be precise we need data on the predicted loss of species owing to each threat. For example, because our estimates are driven largely by changes in human population, ecoregions with expected population declines are predicted to lose no species, regardless of conservation activity. Thus, in

our Mediterranean habitat case study, we predicted there would be zero loss of species in the Mountain Fynbos and Lowland Fynbos of South Africa due to anticipated human population declines in these areas. A more detailed analysis of local threats would likely yield a different conclusion. Lastly, when proposing a conservation action, rarely is there any discussion of the real possibility of failure, which should be used to discount the return on the investment. All of these challenges for ROI are also challenges for more conventional resource allocation approaches – they are simply more transparent when formally calculating a ROI.

ROI is not just about giving answers – it provides a structured way of learning. For example, the ROI approach allows us to carry out sensitivity analyses and pinpoint which additional data we need most. This learning property is in contrast to alternative approaches, such as the use of expert panels. While expert opinion is useful, the only way to improve it is to get "better" experts. Experts tell us what they know; ROI analyses can tell us we need information in areas the experts might not have considered. Finally, ROI analyses are transparent, whereas experts reach conclusions via an internal decision-making process. Indeed, by being explicit, transparent and quantitative, the ROI approach forces us to specify factors that are usually at best incompletely acknowledged.

We focused here on the ecoregion as the conservation unit, but could have used smaller (or larger) spatial scales, for example allocating resources among sites within or across ecoregions. In other words, ROI analyses are completely scalable. Of course analyses at a different scale are likely to yield a different blend of recommended allocations. Our analyses effectively assume each ecoregion is homogeneous with respect to species richness, costs, etc. This assumption is certainly false. Some sites in an ecoregion will be unusually species-rich, others unusually species-poor, and costs will vary in space, so we expect different allocations if we analyze at different spatial scales. Analyses at smaller scales can consider explicitly issues such as habitat-requirements for different species, connectivity of parcels, the need for multiple sites, etc.

Most conservation practitioners are aware of most or all of the issues discussed in this paper. They know costs are crucial, that urgency needs to be traded off against availability, and so on. They also are aware of many local factors – e.g. landowner sentiment and connections between sites and funding opportunities – that may be complicated and possibly (though not necessarily) hard to quantify. Practitioners factor this knowledge, informed by their experience, into their decisions. Nevertheless, while mentally trading off various factors may lead to good decisions, this is a fallible process, opaque in its operation, and ultimately not capable of computing benefit to cost ratios affected by many interdependent components.

This is not to say that ROI is a stand-alone tool for reaching decisions. Obviously, the ROI approach will not provide perfect solutions nor, in many cases, ready-to-use answers. But it can and should inform decisions. And where the final decision is different from that suggested by ROI, clarity will be gained by an explicit statement of the reasons.

We view ROI as a natural next step in the evolution of conservation planning and priority-setting. Twenty years ago formal tools for conservation planning did not exist, and land acquisition was often opportunistic, albeit valuable. Now conservation groups around the world use conservation planning tools and rigorous site-selection algorithms (Sarkar et al., 2006). All of these improvements in planning are motivated by a desire to be more effective and to use our money and effort more wisely, goals for which ROI is precisely designed.

Acknowledgments

This research was funded in part by The University of Queensland, the Australian Research Council, The Nature Conservancy of California, The Oracle Corporation, and The Rodney Johnson/Katherine Ordway Stewardship Endowment and the University of California, Santa Barbara. David Bael and Michael Bode commented on the ms and helped with analyses. Early collaborations were supported by the National Center for Ecological Analysis and Synthesis, University of California, Santa Barbara.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.biocon.2007.07.011.

REFERENCES

- Ando, A., Camm, J., Polasky, S., Solow, A., 1998. Species distributions, land values, and efficient conservation. Science 279, 2126–2128.
- Armsworth, Daily, P.G.C., Kareiva, P., Sanchirico, J.N., 2006. Land market feedbacks can undermine biodiversity conservation. Proceedings of the National Academy of Sciences 103, 5403–5408.
- Balmford, A., Gaston, K.J., Rodrigues, A.S.L., James, A., 2000. Integrating costs of conservation into international priority setting. Conservation Biology 14, 1–9.
- Brooks, T., Mittermeier, R., da Fonseca, G., Gerlach, J., Hoffman, M., Lamoreux, J., Mittermeier, C., Pilgrim, J., Rodriguez, A., 2006. Global biodiversity conservation priorities. Science 313, 58–61.
- Cabeza, M., Moilanen, A., 2003. Site-selection algorithms and habitat loss. Conservation Biology 17, 1402–1413.
- Camm, J.D., Norman, S.K., Polasky, S., Solow, A.R., 2002. Nature reserve selection to maximize expected species coverage. Operations Research 50 (6), 946–955.
- Church, R.L., Stoms, D.M., Davis, F.W., 1996. Reserve selection as a maximal coverage problem. Biological Conservation 76 (2), 105–112.
- Costello, C., Polasky, S., 2004. Dynamic reserve site selection. Resource and Energy Economics 26, 157–174.
- Cowling, R.M., Rundel, P.W., Lamont, B.B., Arroyo, M.K., Arianoutsou, M., 1996. Plant diversity in Mediterraneanclimate regions. Trends in Ecology and Evolution 11, 362–366.
- Csuti, B., Polasky, S., Williams, P.H., Pressey, R.L., Camm, J.D., Kershaw, M., Kiester, A.R., Downs, B., Hamilton, R., Huso, M., Sahr, K., 1997. A comparison of reserve selection algorithms using data on terrestrial vertebrates in Oregon. Biological Conservation 80, 83–97.
- Ferraro, P.J., 2003. Assigning priority to environmental policy interventions in a heterogeneous world. Journal of Policy Analysis and Management 22, 27–43.

Ginn, W.J., 2005. Investing in Nature: Case Studies of Land Conservation in Collaboration with Business. Island Press, Washington DC, 209 pp.

Groves, C.R., Jensen, D.B., Valutis, L.L., Redford, K.H., Shaffer, M.L., Scott, J.M., Baumgartner, J.V., Higgins, J.V., Beck, M.W., Anderson, M.G., 2002. Planning for biodiversity conservation: Putting conservation science into practice. Bioscience 52, 499–512.

Hopper, S.D., Gioia, P., 2004. The Southwest Australian Floristic Region: evolution and conservation of a global hot spot of biodiversity. Annual Review of Ecology and Evolution Systematics 35, 623–650.

Keeney, R.L., Raiffa, H., 1976. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. John Wiley & Sons, New York.

Kier, G., Mutke, J., Dinerstein, E., Ricketts, T.H., Küper, W., Kreft, H., Kier, G., Mutke, J., Dinerstein, E., Ricketts, T.H., Küper, W., Kreft, H., Barthlott, W., 2005. Global patterns of plant diversity and floristic knowledge. Journal of Biogeography 32, 1107–1116.

Mace, G.M., Possingham, H.P., Leader-Williams, N., 2006. Prioritizing choices in conservation. In: MacDonald, D. (Ed.), Key Topics in Conservation Biology. Oxford University Press, UK, pp. 17–34.

Meir, E., Andelman, S., Possingham, H.P., 2004. Does conservation planning matter in a dynamic and uncertain world? Ecology Letters 7, 615–622.

Millennium Ecosystem Assessment (MEA), 2005. Ecosystems and Human Well-being: Policy Responses. Island Press, Washington, DC.

Moilanen, A., Franco, A.M.A., Early, R.I., Fox, R., White, B., Thomas, C.D., 2005. Prioritizing multiple-use landscapes for conservation: methods for large multi-species planning problems. Proceedings of the Royal Society, Series B 272, 1885–1891.

Moilanen, A., Wintle, B.A., Elith, J., Burgman, M., 2006. Uncertainty analysis for regional-scale reserve selection. Conservation Biology 20, 1688–1697.

Moore, J., Balmford, A., Allnutt, T., Burgess, N., 2004. Integrating costs into conservation planning across Africa. Biological Conservation 117, 343–350.

Myers, N., Mittermeier, R.A., Mittermeier, C.G., da Fonseca, G.A.B., Kent, J., 2000. Biodiversity hotspots for conservation priorities. Nature 403, 853–858.

Naidoo, R., Adamowicz, W.L., 2005. Economic benefits of biodiversity conservation exceed costs of conservation at an African rainforest reserve. Proceedings of the National Academy of Sciences of the United States of America 102, 16712–16716.

Naidoo, R., Balmford, A., Ferraro, P.J., Polasky, S., Ricketts, T.H., Rouget, M., 2006. Integrating economic cost into conservation planning. Trends in Ecology and Evolution 21 (12), 681–687.

Nicholson, E., Westphal, M.I., Frank, K., Rochester, W.A., Pressey, R.L., Lindenmayer, D.B., Possingham, H.P., 2006. A new method for conservation planning for the persistence of multiple species. Ecology Letters 9 (9), 1049–1060. Polasky, S., Camm, J., Garber-Yonts, B., 2001. Selecting biological reserves cost-effectively: an application to terrestrial vertebrate conservation in Oregon. Land Economics 77 (1), 68–78.

Polasky, S., Camm, J.D., Slow, A.R., Csuti, B., White, D., Ding, R., 2000. Choosing reserve networks with incomplete species information. Biological Conservation 94, 1–10.

Polasky, S., Nelson, E., Lonsdorf, E., Fackler, P., Starfield, A., 2005. Conserving species in a working landscape: land use with biological and economic objectives. Ecological Applications 15 (4), 1387–1401.

Possingham, H., Ball, I., Andelman, S., 2000. Mathematical methods for identifying representative reserve networks. In: Ferson, S., Burgman, M. (Eds.), Quantitative Methods for Conservation Biology. Springer-Verlag, New York, pp. 291–305.

Possingham, H.P., Andelman, S.J., Noon, B.R., Trombulak, S., Pulliam, H.R., 2001. Making smart conservation decisions. In: Orians, G., Soule, M. (Eds.), Research Priorities for Conservation Biology. Island Press, Washington, DC, pp. 225–234.

Rissman, A., Lozier, L., Comendant, T., Kareiva, P., Kiesecker, J., Shaw, R., Merenlender, A., 2007. Conservation easements: biodiversity protection and private use. Conservation Biology 21, 709–718.

Sanderson, E.W., Jaiteh, M., Levy, M.A., Redford, K.H., Wannebo, A.V., Woolmer, G., 2002. The human footprint and the last of the wild. BioScience 52, 891–904.

Sarkar, S., Pressey, R.L., Faith, D.P., Margules, C.R., Fuller, T., Stoms, D.M., Moffett, A., Wilson, K.A., Williams, K.J., Williams, P.H., Andelman, S., 2006. Biodiversity conservation planning tools: present status and challenges for the future. Annual Review of Environment and Resources 31, 123–159.

Schoener, T.W., 1971. Theory of feeding strategies. Annual Review of Ecology and Systematics 2, 369–404.

Underwood, E.C., Shaw, R., Wilson, K.A., Kareiva, P.K., Klausmeyer, K.R., McBride, M.F., Bode, M., Morrison, S.A., Hoekstra, J.M., Possingham, H.P., Protecting biodiversity when money matters: maximizing return on investment, unpublished ms.

Wilson, K.A., Westphal, M.I., Possingham, H.P., Elith, J., 2005. Sensitivity of conservation planning to different approaches to using predicted species distribution data. Biological Conservation 122, 99–112.

Wilson, K.A., McBride, M.F., Bode, M., Possingham, H.P., 2006. Prioritising global conservation efforts. Nature 440, 337–340.

Wilson, K.A., Underwood, E.C., Morrison, S.A., Klausmeyer, K.R., Murdoch, W.W., Reyers, B., Wardell-Johnson, G., Marquet, P.A., Rundel, P.W., McBride, M.F., Pressey, R.L., Bode, M., Hoekstra, J.M., Andelman, S.J., Looker, M., Rondinini, C., Kareiva, P., Shaw, M.R., Possingham, H.P., 2007. Maximising the conservation of the world's biodiversity: What to do, where and when. PLoS Biology 5, e223.

World Wildlife Fund. 2005. WildFinder: Online database of species distributions, ver 1.0. Available from: <www.worldwildlife.org/ wildfinder>.