Author/s:
Moore, AL;Walker, L;Runge, MC;McDonald-Maden, E;McCarthy, MA

Title:
Two-step adaptive management for choosing between two management actions

Date:
2017-06-01

Citation:
Moore, A. L., Walker, L., Runge, M. C., McDonald-Maden, E. & McCarthy, M. A. (2017). Two-step adaptive management for choosing between two management actions. Ecological Applications, 27 (4), pp.1210-1222. https://doi.org/10.1002/eap.1515.

Persistent Link:
https://hdl.handle.net/11343/292380

Received Date: 07-Jun-2016

Revised Date: 04-Nov-2016

Accepted Date: 09-Dec-2016

Article Type: Articles

Final version received 24 January 2017

Running head: Two-step adaptive management

Two-step Adaptive Management for choosing between two management actions

Authors:

Alana L. Moore[1] (*corresponding author*), School of Biosciences, The University of Melbourne, Parkville VIC 3010 Australia; Unité de Mathématiques et Informatique Appliquées (MIAT), Toulouse INRA, Auzeville BP 52627 31326 cedex France.

Leila Walker[2], RSPB Centre for Conservation Science, RSPB, The Lodge, Sandy, Bedfordshire, SG19 2DL UK.

Michael C. Runge[3] , United States Geological Survey, Patuxent Wildlife Research Centre, 12100 Beech Forest Road, Laurel, MD 20708 USA.

Eve McDonald-Madden[4], Centre for Biodiversity and Conservation Science, School of Geography, Planning and Environmental Management, University of Queensland, St Lucia QLD 4069 Australia.

Michael A. McCarthy[5], School of Biosciences, The University of Melbourne, Parkville VIC 3010 Australia

[1] moa@unimelb.edu.au; [2]leilawalker@hotmail.com; [3]mrunge@usgs.gov;
[3]e.mcdonaldmadden@uq.edu.au; [5]mamcca@unimelb.edu.au

ABSTRACT

Adaptive management is widely advocated to improve environmental management. Derivations of optimal strategies for adaptive management, however, tend to be case specific and time consuming. In contrast, managers might seek relatively simple guidance, such as insight into when a new potential management action should be considered, and how much effort should be expended on trialing such an action.

We constructed a two time-step scenario where a manager is choosing between two possible management actions. The manager has a total budget which can be split between a learning phase and an implementation phase. We use this scenario to investigate when and how much a manager should invest in learning about the management actions available. The optimal investment in learning can be understood intuitively by accounting for the expected value of sample information, the benefits that accrue during learning, the direct costs of learning, and the opportunity costs of learning.

We find that the optimal proportion of the budget to spend on learning is characterized by several critical thresholds that mark a jump from spending a large proportion of the budget on learning to spending nothing. For example, as sampling variance increases it is optimal to spend a larger proportion of the budget on learning, up to a point - if the sampling variance passes a critical threshold, it is no longer beneficial to invest in learning. Similar thresholds are observed as a function of the total budget and the difference in the expected performance of the two actions. We illustrate how this model can be applied using a case study of choosing between alternative rearing diets for hihi, an endangered New Zealand passerine.

Although the model presented is a simplified scenario, we believe it is relevant to many management situations. Managers often have relatively short time horizons for management, and might be reluctant to consider further investment in learning and monitoring beyond collecting data from a single time period.

KEYWORDS: adaptive management, decision analysis, monitoring costs, expected value of perfect and sample information, Bayesian experimental design.

INTRODUCTION

Adaptive management is widely advocated to improve environmental management, and to help determine appropriate levels of monitoring effort to support better management decisions (Walters and Hilborn 1978, Walters and Holling 1990, Johnson and Williams 2015). Adaptive management aims to strike a balance between learning about the system being managed, and actually managing it (Holling 1978, Walters 1986), a balance referred to as the "dual-control problem" in the literature on operations research (Wittenmark 1995). Learning about a system entails both monitoring costs and lost opportunity costs, since experiments in which two or more actions are trialed concurrently inevitably means that a sub-optimal action will be at least partly implemented. Thus, learning about the system will draw on resources that might be used for management. However, the information gained from monitoring and experimentation might improve management in the future. Adaptive management aims to balance the longer-term benefits of learning with its shorter-term costs, helping to determine the appropriate investment in learning.

The academic literature on adaptive management has proliferated, yet examples of successful implementation are rare (Johnson and Williams 2015). Various reasons restrict the use of adaptive management including lack of institutional support and commitment, and insufficient funding for adequate monitoring programs (Walters 2007, Johnson and Williams 2015). The computational burden required to optimize adaptive management is another potential concern (Martell and Walters 2008). Further, the academic literature tends to emphasize solutions to specific adaptive management problems (e.g. Gregory et al. 2006, Tyre et al. 2011, Shea et al. 2014) ,and drawing general conclusions appears difficult. In contrast, managers might seek relatively simple guidance, such as insight into when a new potential management action should be considered, and how much effort should be expended on trialing such actions (Walters and Green 1997, McDonald-Madden et al. 2010).

To help generalize adaptive management beyond individual case studies, we constructed an adaptive management problem where a manager is choosing between two possible management actions in two decision phases – a learning phase and a final decision phase. The manager has a total budget to spend over these two phases. In the first time period, both actions can be implemented and the results monitored. At the end of this learning phase, the remaining budget will be spent on implementing the action with the highest expected efficiency. The management goal is to maximize the total expected benefit over the two phases. We use this framework to investigate the following questions. First, when should we invest in learning more about the value of the management actions? Second, if investing in a learning phase is expected to be beneficial, how much of the total budget should we invest? Third, how should the amount spent on the learning phase be split between the two available actions given their current expected performance and our uncertainty about these values? Finally, when do we expect the largest benefits from investing in learning?

While this is a simplified scenario, we believe it is relevant to many management situations (e.g. see the frameworks proposed by Walters and Green 1997, MacGregor et al. 2002). The two phases will at best approximate sequential decisions over many time-steps, however, managers often have relatively short time horizons for management, and might be reluctant to consider further investment in experimentation and monitoring beyond collecting data from a single time period. As we show in this paper, one advantage of this simplified scenario is that analytical expressions for the optimal level of experimentation (i.e. optimal number of samples of each management action during the first phase) can be obtained for particular special cases, and numerical solutions can be obtained efficiently in other cases.

There exists a substantial literature addressing the optimal sample size when choosing between two, or more, treatments in clinical trials, when the objective is to maximize the total number of successful treatments (Hardwick and Stout 2002, Ghosh et al. 2011). However, we found that these studies consider scenarios that differ from that considered here in one or more of the following four respects: (i) the cost of performing experiments is ignored (e.g. Colton 1963), (ii) the number of trials is assumed to be the same for the two treatments (e.g. Canner 1970, Willan and Kowgier 2008), (iii) they consider dichotomous responses (success or failure) from the trials (e.g.Cheng 1996, Hardwick and Stout 2002, Cheng et al. 2003), or (iv) they consider testing a new action against a known one (e.g. Grundy et al. 1954). As far as we are aware, the scenario considered in this study (including sample costs, unequal allocation of trials during the experimental phase, a measure of benefit size obtained from each trial, and two uncertain actions) has not been addressed in this literature nor in the literature on natural resource management.

Another approach to evaluating the expected value of experimentation is value of information (VOI) analysis (Raiffa and Schlaiffer 1961). VOI is a broad term for an analysis that estimates the expected potential value of gaining new information about a system. VOI has been used in various disciplines to determine the maximum amount that should be invested in gaining information before making a decision (Maxwell et al. 2015). In particular, VOI has been applied to environmental management dilemmas to determine the potential management benefit of resolving uncertainty both for one off (Runge et al. 2011, Maxwell et al. 2015) and dynamic decision processes (Williams et al. 2011, Williams and Johnson 2015), and to determine whether or not monitoring should be performed (Hauser et al. 2006, McDonald-Madden et al. 2010).

VOI analyses may consider the value of resolving all uncertainty about a system (Expected Value of Perfect Information, EVPI), the value of resolving some sources of uncertainty (expected value of partial information), or the value of resolving some of the uncertainty via additional sampling (expected value of sample information, EVSI) (Runge et al. 2011). Such analyses provide an upper bound on how much should be invested in gathering information before taking a management decision, and can identify when the benefits of learning are expected to be the greatest. However, such analysis does not tell us the optimal amount to invest in learning when accounting for monitoring and lost opportunity costs.

At least two decision phases must be considered to capture the trade-off between the expected benefits and costs of experimentation. We relate the solution of our two time-step process to the EVSI, and highlight the trade-off between the value of sample information and lost opportunity costs. By nesting the experimental design question within a decision question, we take the same approach as in Bayesian experimental design (Chaloner and Verdinelli 1995); indeed, EVSI is very closely related to a Bayesian preposterior analysis, and provides similarly relevant information to a decision maker.

METHODS

We consider the case when a manager has two actions to choose from, $i = \{1, 2\}$. The manager has a total budget B to spend on implementing the actions. For each action, one unit of management costs c_i and results in a benefit x_i . We assume that the benefit of each action is uncertain such that x_i is an unknown random variable, with the uncertainty represented by a normal distribution with mean m_i and standard deviation s_i.

Before presenting the two-step adaptive management model, we first consider the expected value of sample information. This gives us the expected benefit of information acquired from a particular experimental design. EVSI essentially ignores costs associated with obtaining the experimental results; whether or not experimentation occurs, the same amount will be invested in implementing the expected best action. We then consider a two-step adaptive management (AM) scenario, made up of an experimental phase and an implementation phase. We use this framework to investigate the trade-off between investing in experimentation and saving resources to implement the best action. We highlight the relationship between the AM solution and EVSI.

Expected Value of Sample Information

In the case that the manager must choose between the two actions in the absence of any further information, or reduction in uncertainty, the optimal decision is to invest the entire budget in the action i that maximizes the expected net benefit, with the expectation taken over the prior distribution. The expected net benefit in the face of uncertainty is

$$E_u = \max_i E\left[B\frac{x_i}{c_i}\right] = B\max\left(\frac{m_1}{c_1},\frac{m_2}{c_2}\right) = B\left\{\frac{m_1}{2c_1}+\frac{m_2}{2c_2}+\frac{1}{2}\left|\frac{m_1}{c_1}-\frac{m_2}{c_2}\right|\right\}. \qquad (1)$$

The expected value of sample information (EVSI) is the difference between the expected value after a given sampling regime is implemented (reduction but not elimination of uncertainty) and the expected value in the face of uncertainty. Hence, to calculate EVSI, we need to calculate the pre-posterior distribution, that is, the expected net benefit from having additional information, taken with respect to the prior distribution.

Suppose that our sampling design is to observe n_1 units of action 1 and n_2 units of action 2. We then observe a mean response p_i for the units under action i, and these have an individual variation of σ_i^2. We assume the p_i are independently distributed according to

$$p_i | x_i \sim N\left(x_i, \sqrt{\sigma_i^2 / n_i}\right) \quad (2)$$

and the unconditional distribution, given the prior for x_i, is

$$p_i \sim N\left(m_i, \sqrt{\sigma_i^2 / n_i + s_i^2}\right). \quad (3)$$

Combining the prior and the observed data, using Bayes' Theorem, the posterior distribution for the per-unit benefit, y_i, is normal with mean

$$m_i' = \frac{p_i n_i s_i^2 + m_i \sigma_i^2}{n_i s_i^2 + \sigma_i^2} \quad (4)$$

and variance

$$\sigma_i'^2 = \frac{\sigma_i^2 s_i^2}{n_i s_i^2 + \sigma_i^2}. \quad (5)$$

After observing the new information, we would choose the action with the highest expected efficiency, with the expectation taken over the posterior distribution,

$$\max\left(E_{posterior}\left[\frac{y_1}{c_1}, \frac{y_2}{c_2}\right]\right) = max\left(\frac{p_1 n_1 s_1^2 + m_1 \sigma_1^2}{c_1(n_1 s_1^2 + \sigma_1^2)}, \frac{p_2 n_2 s_2^2 + m_2 \sigma_2^2}{c_2(n_2 s_2^2 + \sigma_2^2)}\right)$$

$$= \frac{1}{2}\left(\frac{p_1 n_1 s_1^2 + m_1 \sigma_1^2}{c_1(n_1 s_1^2 + \sigma_1^2)} + \frac{p_2 n_2 s_2^2 + m_2 \sigma_2^2}{c_2(n_2 s_2^2 + \sigma_2^2)} + \left|\frac{p_1 n_1 s_1^2 + m_1 \sigma_1^2}{c_1(n_1 s_1^2 + \sigma_1^2)} - \frac{p_2 n_2 s_2^2 + m_2 \sigma_2^2}{c_2(n_2 s_2^2 + \sigma_2^2)}\right|\right). \quad (6)$$

Because we wish to estimate this value prior to making the observation of $\{p_1, p_2\}$, we now need to take the expectation of this quantity with respect to the prior distribution. The only random variables are p_1 and p_2. Thus, the pre-posterior expectation for the maximum efficiency is

$$E_e = \frac{1}{2}\left(\frac{E[p_1] n_1 s_1^2 + m_1 \sigma_1^2}{c_1(n_1 s_1^2 + \sigma_1^2)} + \frac{E[p_2] n_2 s_2^2 + m_2 \sigma_2^2}{c_2(n_2 s_2^2 + \sigma_2^2)} + E[|\Delta|]\right) = \frac{1}{2}\left(\frac{m_1}{c_1} + \frac{m_2}{c_2} + E[|\Delta|]\right) \quad (7)$$

where

$$\Delta = \frac{p_1 n_1 s_1^2 + m_1 \sigma_1^2}{c_1(n_1 s_1^2 + \sigma_1^2)} - \frac{p_2 n_2 s_2^2 + m_2 \sigma_2^2}{c_2(n_2 s_2^2 + \sigma_2^2)} \quad (8)$$

is normally distributed with mean

$$\mu = \frac{m_1}{c_1} - \frac{m_2}{c_2} \tag{9}$$

and variance

$$\Theta^2 = \left(\frac{s_1}{c_1}\right)^2 \frac{n_1 s_1^2}{n_1 s_1^2 + \sigma_1^2} + \left(\frac{s_2}{c_2}\right)^2 \frac{n_2 s_2^2}{n_2 s_2^2 + \sigma_2^2}. \tag{10}$$

Because Δ is normally distributed, the modulus (absolute value) of Δ has a folded-normal distribution. Thus,

$$E_e = \frac{1}{2}\left(\frac{m_1}{c_1} + \frac{m_2}{c_2} + \Theta\sqrt{\frac{2}{\pi}}\, e^{-\mu^2/2\Theta^2} + \mu \operatorname{erf}\left(\frac{\mu}{\Theta\sqrt{2}}\right)\right). \tag{11}$$

In the case that sampling is obtained for free and the entire budget *B* is spent on implementing the action with the highest expected posterior efficiency, the total expected benefit with sampling is

$$E_s = B * E_e,$$

$$= \frac{B}{2}\left(\frac{m_1}{c_1} + \frac{m_2}{c_2} + \Theta\sqrt{\frac{2}{\pi}}\, e^{-\mu^2/2\Theta^2} + \mu \operatorname{erf}\left(\frac{\mu}{\Theta\sqrt{2}}\right)\right). \tag{12}$$

The expected value of sample information (EVSI) is the difference between the expected benefit with sampling and the expected benefit in the face of uncertainty:

$$EVSI = E_s - E_u,$$

$$= \frac{B}{2}\left\{\frac{m_1}{c_1} + \frac{m_2}{c_2} + \Theta\sqrt{\frac{2}{\pi}}\, e^{-\mu^2/2\Theta^2} + \mu \operatorname{erf}\left(\frac{\mu}{\Theta\sqrt{2}}\right)\right\} - \frac{B}{2}\left\{\frac{m_1}{c_1} + \frac{m_2}{c_2} + \left|\frac{m_1}{c_1} - \frac{m_2}{c_2}\right|\right\},$$

$$= \frac{B}{2}\left\{\Theta\sqrt{\frac{2}{\pi}}e^{-\mu^2/2\Theta^2} + \mu\,\mathrm{erf}\left(\frac{\mu}{\Theta\sqrt{2}}\right) - |\mu|\right\}, \quad (13)$$

where $\mu = \frac{m_1}{c_1} - \frac{m_2}{c_2}$, and $\Theta^2 = \left(\frac{s_1}{c_1}\right)^2 \frac{n_1 s_1^2}{n_1 s_1^2+\sigma_1^2} + \left(\frac{s_2}{c_2}\right)^2 \frac{n_2 s_2^2}{n_2 s_2^2+\sigma_2^2}$ (Eqs. 9 and 10). The derivation of the expected value of perfect information (EVPI) and a comparison with EVSI can be found in Appendix S1.

Two-step Adaptive Management

To calculate the EVSI we assumed that we knew the sampling design; the number of units n_i of each action to be trialed. The EVSI tells us the maximum *additional* amount we could spend on a particular sampling design to achieve the same expected net benefit. However, in the case that we have a total budget B to spend on both experimentation and implementation, EVSI does not tell us how much of that budget to invest in monitored trials. The more we invest in experimentation, the more likely we are to finally choose the best management action, but experimentation incurs additional monitoring costs (resulting in less money to spend on implementation) and lost opportunity costs of trialing the worst action.

To analyze this trade-off we consider a two-step adaptive management process, made up of an experimental phase, consisting of monitored trials, and an implementation phase, in which the remaining funds are used to implement the action with the largest posterior efficiency. During the experimental phase, the additional cost of monitoring the outcome of each trial is k_i for each unit of management. We assume that the data will have a standard deviation of σ_i, representing the observed variation in benefit among different units of management.

The total expected net benefit over the two time-steps is the expected benefit from the experimental phase plus the expected benefit of spending the remaining funds on the action that is found to have the highest expected efficiency (equation 11),

$$L = n_1 m_1 + n_2 m_2 + (B - n_1(c_1 + k_1) - n_2(c_2 + k_2))E_e,$$

$$= n_1 m_1 + n_2 m_2 + \left(B - n_1(c_1+k_1) - n_2(c_2+k_2)\right)\frac{1}{2}\left\{\frac{m_1}{c_1} + \frac{m_2}{c_2} + \Theta\sqrt{\frac{2}{\pi}}e^{-\mu^2/2\Theta^2} + \mu \operatorname{erf}\left(\frac{\mu}{\Theta\sqrt{2}}\right)\right\},$$

(14)

Let the total cost of the experimental phase be given by $C_{experiment} = (c_1+k_1)n_1 + (c_2+k_2)n_2$. Equation (14) can be re-written as

$$L = n_1 m_1 + n_2 m_2 + \left(1 - \frac{C_{experiment}}{B}\right) E_s$$

$$= n_1 m_1 + n_2 m_2 + E_s - \frac{C_{experiment}}{B} E_s$$

$$= n_1 m_1 + n_2 m_2 + E_u + EVSI - \frac{C_{experiment}}{B} E_s$$

$$= E_u + EVSI + n_1 m_1 + n_2 m_2 - \frac{C_{experiment}}{B} E_s. \qquad (15)$$

Written in this way we can more easily see the trade-off between investing in experimentation and saving resources for implementing the best action; the more we spend on the experimental phase, the larger the expected value of sample information (EVSI) and the larger the incidental benefits of experimentation ($n_1 m_1 + n_2 m_2$). However, the lost opportunity costs incurred by using up resources during sampling, $(C_{experiment}/B)E_s$, are also greater.

The number of trials of each action that maximizes the total net benefit can be found efficiently using numerical methods. We generated the numerical results using Wolfram Mathematica V8.0.4 (Inc. 2014)). We used a built in optimization function, FindMaximum, to find the optimal non-zero allocation to the learning phase and compared this to the expected reward under no experimentation (Data S1). We also derived explicit analytic solutions for several special cases (Appendix S2).

Example: Choosing between supplementary feeding options for hihi nestlings

We illustrate the model by determining the optimal proportion of the total budget to use on trialing two supplemental feeding treatments for hihi (*Notiomystis cinta*) nestlings, an endangered New Zealand bird whose recovery program is based on supplementary feeding (Walker et al. 2013). There is evidence that sugar water improves adult survival (Armstrong and Ewan 2001;

Chauvenet et al 2012); sugar water is currently provided to five out of six extant populations (L.Walker, personal observations). An alternative full dietary supplement (WombarooTM Lorikeet & Honeyeater Food, Wombaroo Food Products, Glen Osmond, SA, Australia) has also been trialed in both adult and, more recently, juvenile populations (Armstrong et al. 2007, Walker et al. 2013). Walker et al. (2013) investigated experimentally the effects of neonatal supplementary feeding using four alternative treatments on nestling growth, nestling survival and juvenile survival to breeding age (recruitment). The following illustrative example is based on data and cost estimates from Walker et al.'s study.

Consider the case when management has a total budget *B* to spend on supplementary feeding over *T* years. The manager has two possible supplementary feeding treatments: sugar water (N-) and WombarooTM Lorikeet & Honeyeater Food (N+). The goal is to determine the proportion of the budget to spend on trialing the two treatments in the first year. The management benefit of each treatment is measured as the mean additional weight at age 20 days; where additional is in reference to the expected average weight with no supplementary feeding. We consider the management units to be birds and consider costs in units of hours per bird per year.

We assume that sugar water is provided using general feeding stations in all situations (i.e. during the experiment and during the management only phase). During the experimental phase, the managers additionally feed the dietary supplement to the nestlings directly. If sugar water (N-) is found to be the preferable treatment, then it would be administered only via the general feeding stations, as it is known to be provisioned to nestlings by parents (Walker et al. 2013; Thorogood et al. 2008). However, for Wombaroo (N+), it is unclear whether it would be possible to administer the supplement via the feeders or if it would be necessary to continue directly feeding juveniles in the nests (L. Walker, personal observations). Therefore, we considered two scenarios. In scenario (i) we assumed that the full dietary supplement (N+) will continue to be administered to juveniles directly. In scenario (ii) we assumed that after the experimental phase N+ could be administered via the general feeding stations. In this case a larger quantity of the dietary supplement would be required, but the cost associated with administering the supplement would be much less.

Estimates for the cost of implementing both management options (general feeders and direct feeding of nestlings), together with estimates of the cost of monitoring the results were obtained from data provided by L. Walker and A. Baxter (unpublished data; personal communication). A summary of the parameters used for the results presented are given in Table 1, while an overview of the cost data can be found in Appendix S3.

RESULTS

When and how much should we invest in learning?

Recall that E_u is the expected benefit in the face of uncertainty, that is, if no experimentation occurs. Consequently, from Equation (15) we see that it is beneficial to invest in experimentation if there exists a sampling design $\{n_1, n_2\}$ (not = $\{0,0\}$) such that the expected benefit from the experimentation phase outweighs the lost opportunity costs incurred by using resources for experimentation, i.e. when

$$EVSI + n_1 m_1 + n_2 m_2 > \frac{C_{experiment}}{B} E_s. \qquad (16)$$

There is no simple rule for when learning is worthwhile due to the large number of parameters involved in determining the threshold. Nonetheless, general tendencies can be observed (summarized in Box 1 and Table 2).

If monitoring costs are negligible it is nearly always optimal to spend some of the budget on learning (Figs. 1-3: Panels a and c, Appendix S4: Fig. S1a). Note that if both actions are uncertain, trialing the expected best action will never be worse than directly implementing it, but there may be no expected advantage when the means are very different. In the case that the benefit of one action is known, if the uncertain action is expected to be worse, then whether or not it is worth trialing it will depend on how uncertain we are about its performance, the budget and the monitoring precision (Figs. 1-3: Panel c, Fig. S2).

If monitoring costs are significant, it is not beneficial to invest in learning if: one action is expected to be much better than the other, monitoring variance is large, monitoring costs are large or the budget is small (Figs. 1-3: Panels b and d). For example, if the expected benefit of the two actions differs, then investing in learning is worthwhile only when the budget is sufficiently large (Fig. 2b and 2d).

Note that the graphs in Figure 1 are not perfectly symmetric around m_2-m_1=0. When the benefit of action 1 is known with certainty (Fig. 1c-d), it is optimal to spend less on the learning phase if the expected benefit of action 2 is less than the expected benefit of action 1 than if it is greater (see also Figs. 2c-d and 3c-d). Intuitively, this is because there is a smaller probability that action 2 is better than action 1. When both actions are uncertain, this argument no longer applies: there is the same probability that the expected worse action will be found to be better. In this case, we observe the opposite behavior: it is optimal to spend a larger proportion of the budget on the learning phase when the expected value of action 2 is 5 units smaller than action 1 than when it is 5 units larger (Figs. 1-3: Panels a-b). This is primarily because the solution depends substantially on the ratio of the means to prior variances: the optimal proportion to spend on the learning phase is a decreasing function of the ratio of the prior expected benefit to prior standard deviation (Appendix S4: Fig. S11).

The solution for the optimal proportion to spend on the learning phase displays a number of interesting critical thresholds (Figs. 1-3). For example, as the difference in the expected prior benefit of the two actions increases, a point is eventually reached beyond which it is not worth investing in learning (Fig. 1). At this point, the optimal solution drops suddenly from spending a large amount on learning to nothing at all. Where this point occurs depends notably on the prior variance of each action. The more uncertain we are about the performance of each action, the greater the difference between the prior mean benefits before we stop investing in learning, since if the overlap between the two prior distributions is small the best action is known with high probability. Similar thresholds are observed for the budget (Fig. 2 and Appendix S4: Fig. S1), monitoring cost (Appendix S4: Fig. S1) and monitoring variance (Fig. 3). These thresholds are more prevalent when monitoring costs are significant.

This threshold behavior can be better understood by observing that the optimal (non-zero) investment in experimentation is a local, but not necessarily global, optimum (Fig. 4). The expected net benefit (ENB = Expected benefit *without* experimentation - expected benefit *with* experimentation) is a concave function of the amount invested in the experiment. Note that no experimentation results in zero expected net benefit. When the expected net benefit of experimentation is positive, the optimal solution is found at the maximum of this curve (e.g., at an investment of ~50 in Fig. 4a). However, as, for example, the sampling variance increases, EVSI and also the expected net benefit decrease, but an optimal allocation can still be found, until the whole curve drops below 0 (Fig. 4b), in which case, no investment in learning is warranted.

Analytical results for the optimal number of trials can be derived for several, potentially common, special cases (Appendix S2). These analytic solutions suggest a maximum of 1/3 of the budget should be spent on the learning phase. Numerical results showed that this limit is occasionally exceeded when: monitoring costs are negligible, means differ, and either sampling variance is (reasonably) high or the budget is small (Figs. 2 and 3). However, for the parameter ranges we explored, it is usually optimal to spend less than 20% of the budget on learning. When monitoring costs are significant, the optimal allocation of effort to the learning phase is always less than a third. Moreover, in this case the analytic solution derived assuming identical parameters (Appendix S2: Eq. S3) is an upper bound.

For both negligible and significant monitoring costs, the highest proportion of the budget is spent on learning when the budget is fairly small (Fig. 2, Appendix S2: Eq. S3). As the budget increases relative to implementation and monitoring costs, we spend more on monitoring in an absolute sense, but a smaller fraction of the total budget. For example, in the hihi supplementary feeding example below, as the budget increases the optimal number of trials of each treatment increases, but the total proportion spent on the learning phase decreases (Fig. 6 and Appendix S4: Fig. S10).

For a fixed budget, the optimal proportion to spend on learning is an increasing function of monitoring costs when parameters are equal and the prior expected benefits are zero (Appendix S2). This is because although the optimal number of trials is a decreasing function of monitoring cost, it does not decrease as fast as monitoring and implementation costs increase. Interestingly, when the prior mean benefits are positive, the optimal number of trials decreases more quickly than when they can be assumed to be zero (Appendix S4: Fig. S3). Consequently, when the management actions are expected to have a large benefit (>~4 for the default parameters), the optimal proportion of the budget to spend on the learning phase is a decreasing, rather than increasing, function of monitoring costs (Appendix S4: Fig. S3).

When monitoring costs are negligible, the amount spent on the learning phase is an increasing function of the difference in the prior mean benefit (until the threshold is reached) (Fig. 1a and 1c). Consequently, the largest percentage of the budget is invested in experimentation when the prior means are different, but not too different. In contrast, when monitoring costs are significant the largest percentage is spent on learning when the prior mean benefits are the same (Fig. 1b and 1d).

When all other parameters are equal, we spend the maximum proportion of the budget on learning when the prior standard deviations of the two actions are the same (Fig. S4). That is, if we are more confident about one action than the other, we will tend to spend less on experimenting. In Appendix S2, we derive an analytic solution for the case when the benefit of one action is well known ($s_1 = 0$), the prior mean benefits are the same and either m_2 or k_2 is zero (Appendix S2: Eq. S6). We also show that this solution is a good approximation for non-zero m_2 and k_2 if the prior variance is large relative to the expected benefit m_2. Our numerical results support this finding: in general, if we fix the uncertainty about one action and increase the uncertainty about the other, the amount spent on learning converges to the analytical solution derived (Appendix S4: Fig. S4). Presumably we are only entertaining the second action because we think that its performance is roughly the same as action 1, but we are not sure whether it will do much better or much worse. In this case, s_2 will be large (relative to m_2). Hence, the analytical result gives us a rough rule of thumb of investment in assessing a very uncertain action against a known outcome.

As highlighted by the analytic results (Appendix S2), the ratio of sampling variance to prior variance also plays an important role in determining when to invest in learning. As the sampling variance increases, more sampling is required to be similarly confident about the benefit of each action, initially increasing the amount spent on the learning phase. However, the percentage gain from investing in learning is a decreasing function of sample variance (Appendix S4: Fig. S5). Consequently, when the expected performance differs between actions or monitoring costs are significant, investing in learning is eventually no longer beneficial. At this critical threshold, the optimal strategy switches from investing a significant amount in learning to investing nothing (Figs. 1 and 3).

If we invest in learning, what is the split between the two actions?

When only the prior mean efficiency differs between actions, it is optimal to spend a larger proportion of the learning-phase budget on the action with the highest expected performance (Fig. 5a-b, s_2 = 10). When the two actions are expected to perform equally well but the uncertainty about their performance differs, it is optimal to spend more on the most uncertain action (Fig. 5a-b, m_2-m_1=0). When the prior mean efficiencies and prior variances both differ, the split is weighted toward the action with the highest expected return (Fig. 5a-b). That is, even if we are more uncertain about action 2, we may still spend more on trialing action 1 if we believe it is expected to be the better action.

When the prior distributions differ, it is sometimes optimal to only trial one of the actions if the budget is small or sampling variance is large, relative to the prior variance, (Fig. 5, Appendix S4: Figs. S6 and S7). In these situations, investing in learning may still be optimal, but lost opportunity costs are minimized by only trialing the expected best action.

When do we get the largest benefits from investing in learning?

The largest percentage gains in the objective function are observed when the budget is large (Appendix S4: Figs. S8b and S9), the means are similar (Appendix S4: Figs. S5, S6a and S9), the efficiency of each action is uncertain (Appendix S4: Fig. S8c), and sampling provides precise results (Appendix S4: Figs. S5 and S8d).

Choosing between supplementary feeding options for hihi nestlings

For the parameters used, the total proportion of the budget to spend on the learning phase depended very little on whether the prior expected benefit of the Wombaroo treatment (N+) was smaller, larger or the same as the sugar water treatment (N-) (Fig. 6a and Appendix S4: Fig. S10a). However, the optimal number of trials of each treatment did depend on the expected benefit of the Wombaroo treatment (Fig. 6b and Appendix S4: Fig. S10b).

Interestingly, there was only a small difference between the results for the two different scenarios (Fig. 6 versus Appendix S4: Fig. S10). That is, for our cost estimates, whether or not Wombaroo would be fed directly to nestlings or could be administered via feeders, the optimal proportion to spend on the learning phase was more or less the same.

It is worth highlighting that these results depend on the reference weight. The optimal proportion to spend on experimentation depends on the ratio of the expected benefit to standard deviation of the prior (Appendix S4: Fig. S11). Consequently, if the reference weight is expected to be large (so that benefit above reference weight is small), then it will be optimal to spend more on experimentation than if the reference weight is small.

Discussion

The formal derivation of the net benefit of two-phase adaptive management for a simple setting provides some powerful intuitive guidance for thinking about the value of learning in a dynamic setting. The value of experimentation arises out of two benefits and two costs (Eq. 15): the benefits associated with applying learning to subsequent management (EVSI); the transient benefits accrued during the learning phase; the direct costs of learning; and the opportunity costs of learning (the resources not available for subsequent management). Experimentation will be warranted when the benefits outweigh the costs (Eq. 16); otherwise, management should proceed in the face of uncertainty. These qualitative insights, derived from quantitative results, provide a useful framework for evaluating experimentation.

More specific guidance for investment in learning becomes complicated quickly (Box 1 and Table 2). The management scenario we have presented was as simple as we could make it while including all the relevant factors. Nevertheless, there were still 11 parameters to consider, making it difficult to extract general insights and tendencies from numerical sensitivity analyses alone. By considering a simple scenario we were able to derive analytical solutions for several special cases. These solutions provided greater insight into how parameter combinations drive solution behavior, and a base against which to compare results when the assumptions leading to an analytical solution are violated. For example, the analytic solution is a good rule of thumb when trialing an unknown management action against a known one, even when the prior expected benefits differ and monitoring costs are significant (contrary to assumptions used to derive the result)(Appendix S4: Fig. S4). However, when considering two uncertain management actions, the optimal allocation of resources depends strongly on the parameters that were excluded from the analytical result.

The scenario analyzed gave rise to several unintuitive results. For example, there is a tendency to think that monitoring large projects is more important than monitoring small projects – sure, large projects should have more money spent on monitoring, but our results suggest that smaller projects should have a higher proportion of the budget spent on learning and monitoring. This also suggests benefits of cooperation and coordination of smaller projects.

An interesting feature of the solution is the existence of several critical thresholds that mark a jump from investing a lot in learning to not learning at all, or vice versa. For example, when the expected performance of the two actions differ, as sampling variance increases we observe a critical threshold at which the optimal solution changes from spending a lot on the learning phase to spending nothing. While it is important for people developing and interpreting adaptive management models to be aware of such thresholds, managers implementing the policies need not be too concerned, since, at these thresholds investing a lot or not investing at all yield quite similar management outcomes. Consequently, it is not crucial to know precisely on which side of the critical threshold the system lies.

The priors for the two management options influence the results quite substantially. That is, the perceived performance of the two options (the prior means), and the uncertainty about their performance (the prior standard deviations) will influence the optimal extent of experimentation. This makes intuitive sense because managers would be expected to entertain the possibility of experimenting on a new management action only if they thought that it might perform better than an alternative but were uncertain about its relative performance. However, prior distributions are rarely used in ecology (Morris et al. 2015), and they can be difficult to specify coherently (McCarthy 2007) . If one were unwilling to specify a prior distribution, then one could set the prior standard deviation to be large, which would mean the posterior distribution would have the same shape as the likelihood function. In this case, the Bayesian estimates of the experimental results would be numerically equivalent to those of a frequentist analysis, which do not incorporate priors. However, such a wide prior distribution implies that extremely good (large positive values for the efficiency of management) or extremely poor outcomes (large negative values) are conceivable. Inflating the uncertainty in the priors will tend to drive more experimentation than might be warranted, emphasizing the need to specify priors thoughtfully with available data (McCarthy and Masters 2005) or rigorous methods for expert elicitation

(Speirs-Bridge et al. 2010). Although priors might be difficult to specify, decision-makers are inherently considering them when they begin to compare different management actions. Explicitly specifying the anticipated benefits and the degree of uncertainty about action outcomes can lead to better decisions about experimentation. Hence, specifying priors should not be seen as an obstacle to the decision making process, but rather a useful tool to improve decisions.

Monitoring is the cornerstone of successful adaptive management (Moir and Block 2001). However, monitoring management outcomes is rarely a trivial task and can account for a large fraction of the total budget required to implement an adaptive approach to management (Walters 2007). For the two time-step process considered here, including monitoring costs substantially changed the solution, both in terms of quantitative value and qualitative behavior. For example, when monitoring costs were negligible, the amount spent on learning increased as the expected benefit of the two actions differed. In this case, little is to be lost by spending more on the learning phase and increasing the proportion of the learning-phase budget spent on the expected best action. However, when monitoring costs were significant the amount spent on experimentation was largest when the difference in the prior mean benefits was small. This is because the resulting probability of choosing the best management action without monitoring is lowest at this point (in contrast, when the expected difference in benefit is large, the probability the better looking action is actually better is large, hence there is less to gain from monitoring, see also MacGregor et al. 2002, Maxwell et al. 2015). Further, the critical thresholds play a more important role when monitoring costs are significant; the minimum budget, maximum difference between prior means and maximum monitoring variance are more likely to be encountered within feasible parameter ranges.

Interestingly, in many ways the optimal solution was simpler when monitoring costs were substantial. For example, the proportion of the budget spent on learning tended to be fairly constant across the region in which it was optimal to invest in learning. Further, the analytic solution derived assuming identical parameters for the two actions (Appendix S2) is an upper-bound on the optimal proportion to spend on learning. These results highlight the importance of accounting for monitoring costs when designing adaptive management plans.

We found that the largest expected proportional gains in the objective function rarely corresponded to when the largest proportion of the budget should be spent on learning. For example, larger proportional gains are expected when sampling variance is low, whereas, in general, a larger percentage of the budget should be spent on learning when sampling variance is high because more samples are needed. Similarly, although a larger percentage gain is expected for large budgets, a larger proportion of the budget should be spent on learning when budgets are small.

We considered a two-step adaptive management approach in which the management horizon is divided into a learning phase and an implementation phase. Walters and Green (1997) propose a similar framework for evaluating experimental management actions for ecological systems. These two approaches make a one-off decision about how much to invest in learning. This differs from many formulations of AM that assume a fixed budget per time-step and look at how to divide funds between alternative management actions, and monitoring, at each phase (e.g. Moore and McCarthy 2010, Baxter and Possingham 2011); effectively deciding how much to invest in learning at each time-step. At best, the two time-steps will approximate sequential decisions over many time-steps. An interesting avenue of future research would be to compare the management policies derived under the two different modelling approaches.

EVSI tells us the expected value of a given sampling design, but it does not take into account lost opportunity costs associated with experimentation and monitoring. Consequently, while methods such as EVSI are useful for determining when learning is likely to be beneficial, and can provide upper bounds on *additional* funds that should be spent on experimentation, further analysis is needed to determine the fraction of the total budget to invest in learning. In contrast, Adaptive Management formulations with long time horizons can be computationally challenging and difficult to implement in the real world. The approach presented here strikes a balance between complexity and utility. By considering a two-step AM process we are able to capture the trade-off between the benefit and costs of investing in additional information while remaining relatively simple.

Acknowledgements

This research was supported by the Australian Research Council (ARC) Centre of Excellence for Environmental Decisions, and an ARC DECRA Award to Eve McDonald-Madden. Any use of

trade, product, or firm names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

References

Baxter, P. W. J. and H. P. Possingham. 2011. Optimizing search strategies for invasive pests: learn before you leap. Journal of Applied Ecology **48**:86-95.

Canner, P. L. 1970. On selecting one of two medical treatments when responses are dichotomous. Journal of the American Statistical Association **65**:293-306.

Chaloner, K. and I. Verdinelli. 1995. Bayesian experimental design: A review. Statistical Science **10**:273-304.

Cheng, Y. 1996. Multistage bandit problems. Journal of Statistical Planning and Inference **53**:153-170.

Cheng, Y., F. S. Su, and D. A. Berry. 2003. Choosing sample size for a clinical trial using decision analysis. Biometrika **90**:923-936.

Colton, T. 1963. A model for selecting 1 of 2 medical treatments. Journal of the American Statistical Association **58**:388-400.

Ghosh, M., N. Mukhopadhyay, and P. K. Sen. 2011. Sequential estimation. John Wiley & Sons.

Gregory, R., L. Failing, and P. Higgins. 2006. Adaptive management and environmental decision making: A case study application to water use planning. Ecological Economics **58**:434-447.

Grundy, P. M., D. H. Rees, and M. J. R. Healy. 1954. Decision between 2 alternatives - how many experiments? Biometrics **10**:317-323.

Hardwick, J. and Q. F. Stout. 2002. Optimal few-stage designs. Journal of Statistical Planning and Inference **104**:121-145.

Hauser, C. E., A. R. Pople, and H. P. Possingham. 2006. Should managed populations be monitored every year? Ecological Applications **16**:807-819.

Holling, C. S. 1978. Adaptive Environmental Assessment and Management. Wiley, London.

Inc., W. R. 2014. Mathematica, Version 8.0.4. Champaign, IL.

Johnson, F. A. and B. K. Williams 2015. A Decision-Analytic Approach to Adaptive Resource Management. Pages 61-84 *in* C. R. A. a. A. S. Garmestani, editor. Adaptive Management of Social-Ecological Systems Springer, Netherlands.

MacGregor, B. W., R. M. Peterman, B. J. Pyper, and M. J. Bradford. 2002. A decision analysis framework for comparing experimental designs of projects to enhance Pacific salmon. North American Journal of Fisheries Management **22**:509-527.

Martell, S. J. D. and C. J. Walters. 2008. Experimental policies for rebuilding depleted stocks. Canadian Journal of Fisheries and Aquatic Sciences **65**:1601-1609.

Maxwell, S. L., J. R. Rhodes, M. C. Runge, H. P. Possingham, C. F. Ng, and E. McDonald-Madden. 2015. How much is new information worth? Evaluating the financial benefit of resolving management uncertainty. Journal of Applied Ecology **52**:12-20.

McCarthy, M. A. 2007. Bayesian Methods for Ecology. Cambridge University Press, Cambridge, UK.

McCarthy, M. A. and P. Masters. 2005. Profiting from prior information in Bayesian analyses of ecological data. Journal of Applied Ecology **42**:1012-1019.

McDonald-Madden, E., P. W. J. Baxter, R. A. Fuller, T. G. Martin, E. T. Game, J. Montambault, and H. P. Possingham. 2010. Monitoring does not always count. Trends in Ecology & Evolution **25**:547-550.

Moir, W. H. and W. M. Block. 2001. Adaptive management on public lands in the United States: Commitment or rhetoric? Environmental Management **28**:141-148.

Moore, A. L. and M. A. McCarthy. 2010. On Valuing Information in Adaptive-Management Models. Conservation Biology **24**:984-993.

Morris, W. K., P. A. Vesk, M. A. McCarthy, S. Bunyavejchewin, and P. J. Baker. 2015. The neglected tool in the Bayesian ecologist's shed: a case study testing informative priors' effect on model accuracy. Ecology and Evolution **5**:102-108.

Runge, M. C., S. J. Converse, and J. E. Lyons. 2011. Which uncertainty? Using expert elicitation and expected value of information to design an adaptive program. Biological Conservation **144**:1214-1223.

Shea, K., M. J. Tildesley, M. C. Runge, C. J. Fonnesbeck, and M. J. Ferrari. 2014. Adaptive Management and the Value of Information: Learning Via Intervention in Epidemiology. Plos Biology **12**.

Speirs-Bridge, A., F. Fidler, M. McBride, L. Flander, G. Cumming, and M. Burgman. 2010. Reducing Overconfidence in the Interval Judgments of Experts. Risk Analysis **30**:512-523.

Tyre, A. J., J. T. Peterson, S. J. Converse, T. Bogich, D. Miller, M. P. van der Burg, C. Thomas, R. Thompson, J. Wood, D. C. Brewer, and M. C. Runge. 2011. Adaptive Management of Bull Trout Populations in the Lemhi Basin. Journal of Fish and Wildlife Management **2**:262-281.

Walters, C. J. 1986. Adaptive Management of Renewable Resources. McGraw Hill, New York.

Walters, C. J. 2007. Is adaptive management helping to solve fisheries problems? Ambio **36**:304-307.

Walters, C. J. and R. Green. 1997. Valuation of experimental management options for ecological systems. Journal of Wildlife Management **61**:987-1006.

Walters, C. J. and R. Hilborn. 1978. Ecological Optimization and Adaptive Management Annual Review of Ecology and Systematics **9**:157-188.

Walters, C. J. and C. S. Holling. 1990. Large-scale management experiments and learning by doing. Ecology **71**:2060-2068.

Willan, A. and M. Kowgier. 2008. Determining optimal sample sizes for multi-stage randomized clinical trials using value of information methods. Clinical Trials **5**:289-300.

Williams, B. K., M. J. Eaton, and D. R. Breininger. 2011. Adaptive resource management and the value of information. Ecological Modelling **222**:3429-3436.

Williams, B. K. and F. A. Johnson. 2015. Value of information in natural resource management: technical developments and application to pink-footed geese. Ecology and Evolution **5**:466-474.

Wittenmark, B. 1995. Adaptive dual control methods: An overview.

Supporting Information

Additional supporting information may be found in the online version of this article at http://onlinelibrary.wiley.com/doi/10.1002/eap.xxxx/suppinfo

Box 1. Summary of key results

When and how much should we invest in learning? (*Figs 1-3, S2 and S4*)

It is worthwhile investing in learning if:

$$EVSI + n_1 m_1 + n_2 m_2 > \frac{c_{experiment}}{B} E_s,$$

for some $\{n_1, n_2\}$ not equal to 0.

Analytic solutions suggest a maximum of 1/3 of the budget should be invested in learning. This is an upper bound when monitoring costs are significant. When monitoring costs are negligible, it may be optimal to spend more than 1/3 on learning if: the prior mean benefit differs between actions and either the sampling variance is (reasonably) large or the budget is small.

The optimal solution is characterized by several interesting critical thresholds.

Significant monitoring costs result in a higher proportion of the budget being spent on learning when the expected performance of the two actions is the same. In contrast, when monitoring costs are negligible a higher proportion of the budget is spent on learning when the expected performance of the two actions differ (Table 2).

How should we split the resources spent on learning between the two actions? (*Fig 5,6 and S6*)

It is optimal to spend more on the most uncertain or the expected best action.

If the prior distributions differ, it is sometimes optimal to only trial one of the actions (expected best or most uncertain) if: the budget is small or sampling variance is large.

Table 1: Parameter estimates for the hihi example.

Parameter	**units**	**Treatment N-**	**Treatment N+**
Management cost, c	hrs/bird/year	(i) 1.13 (ii) 1.13	(i) 2.14 (ii) 2.651
Monitoring cost, k	hrs/bird/year	(i) 1.55 (ii) 1.55	(i) 1.55 (ii) 0.033
Management effect (weight above reference weight) • Mean • SD	g/bird	m_1 = 3 SD = 6	m_2 = {0, 3, 6}g SD = 6
Budget	hours	[50, 2000]	
Monitoring SD/accuracy	g/bird	SD = 6 g	SD = 6 g

Table 2: The largest percentage of the budget is spent on learning when:

variable	**k = 0**	**k > 0**
m_i	m_1 and m_2 *differ* , but difference is < threshold	m_1 and m_2 *are the same*

σ_i	large, but < threshold	
B	small, but > threshold	
s_i	s_1 and s_2 are the same	
k_i	N/A	large if $m_i s_i$ are small small if $m_i s_i$ are large
m_i/s_i	constant	uncertainty about the expected benefit of action *i* is large relative to the expected benefit, i.e when m_i/s_i is small. (Fig S9)

FIGURE LEGENDS

Figure 1. Proportion spent on learning as a function of the difference in the prior expected benefits and sampling standard deviation. Contours indicate the proportion of the budget spent on the learning phase for various shades of gray (dark gray = 0, white = 1). m_1 = 10 (vary m_2), B = 500, $c_1 = c_2 = 5$, $\sigma_1=\sigma_2$. Panels (a) and (c) assume zero monitoring cost, $k_1 = k_2 = 0$, panels (b) and (d) assume a monitoring cost of 3 units ($k_1 = k_2 = 3$). Panels (a) and (b) assume both actions are uncertain with $s_1= s_2 = 10$, panels (c) and (d) assume the benefit of one is known, $s_1 = 0$ and $s_2 = 10$. See Fig S1 for a cross section of (a) and (b) at $\sigma_1 = \sigma_2 = 40$.

Figure 2. Proportion spent on learning vs budget. Default parameters: m_1=10, m_2 = {5, 10, 15}, $s_1 = s_2 = 10$, $c_1=c_2=5$, $\sigma_1 = \sigma_2 = 20$. Black dashed: m_2-m_1=-5, thin black: m_2-m_1=0, thick black: m_2-m_1=5. The gray line is the corresponding analytic solution assuming parameters are equal and either monitoring is free or mean benefits are zero.

Figure 3. Proportion spent on experimentation as a function of the ratio of sample to prior variance for action 2. B=500, m_1=10, $c_1=c_2$=5. (a) & (b) $s_1 = s_2 = 10$, $\sigma_1=\sigma_2$. (c) & (d) benefit of action 1 assumed to be known. Black dashed: m_2=5, thin black: m_2=10, thick black: m_2=15. The

gray line is the corresponding analytic solution assuming parameters are equal and either monitoring is free or mean benefits are zero.

Figure 4. Expected net benefit versus the amount invested in experimentation for: two-step AM = $EVSI + n_1 m_1 + n_2 m_2 - \frac{c_{experiment}}{B} E_s$ (thick solid line), EVSI (dashed line) and EVPI (thin solid line). $c_1 = c_2 = 5$; $k_1 = k_2 = 3$; $s_1 = s_2 = 10$; $m_1 = 10$, $m_2 = 15$; $B = 500$. (a) $\sigma_1 = \sigma_2 = 20$ ($\sigma^2/s^2 = 4$), (b) $\sigma_1 = \sigma_2 = 37$ ($\sigma^2/s^2 = 13.7$).

Figure 5. Panels (a) and (b): Optimal proportion to spend on action 2 when the prior means and standard deviations differ. $s_1 = 10$, $c_1 = c_2 = 5$, $\sigma_1 = \sigma_2 = 20$, $B = 500$. Black dashed: s_2=5, thin black: s_2=10, thick black: s_2=20. A sudden drop to zero corresponds to crossing a threshold above which none of the budget is spent on the learning phase (see panels (c) and (d): optimal proportion to spend on the learning phase). In panel (b), when $s_2 = 5$, none of the budget is spent on trialing action 2 as the benefit is reasonably well known, however, some of the budget is spent on trialing action 1 for some of the parameter space.

Figure 6. Hihi supplementary feeding example, scenario (i). Panel (a): Optimal proportion to spend on the learning phase as a function of the budget (for males). Black dashed: $m_{N+} = 0$, Black thin: $m_{N+} = 3 = m_{N-}$, Black thick: m_{N+}=6. In this panel, the gray line corresponds to the approximate solution, calculate assuming parameters are the same and either negligible monitoring costs or zero expected effect. Panel (b): Corresponding optimal number of trials of each action. Black (top group of lines) = action N- , gray (bottom group) = action N+. Target weight = 29. $s = 6$, $\sigma = 6$. $m_{N-} = 3$.

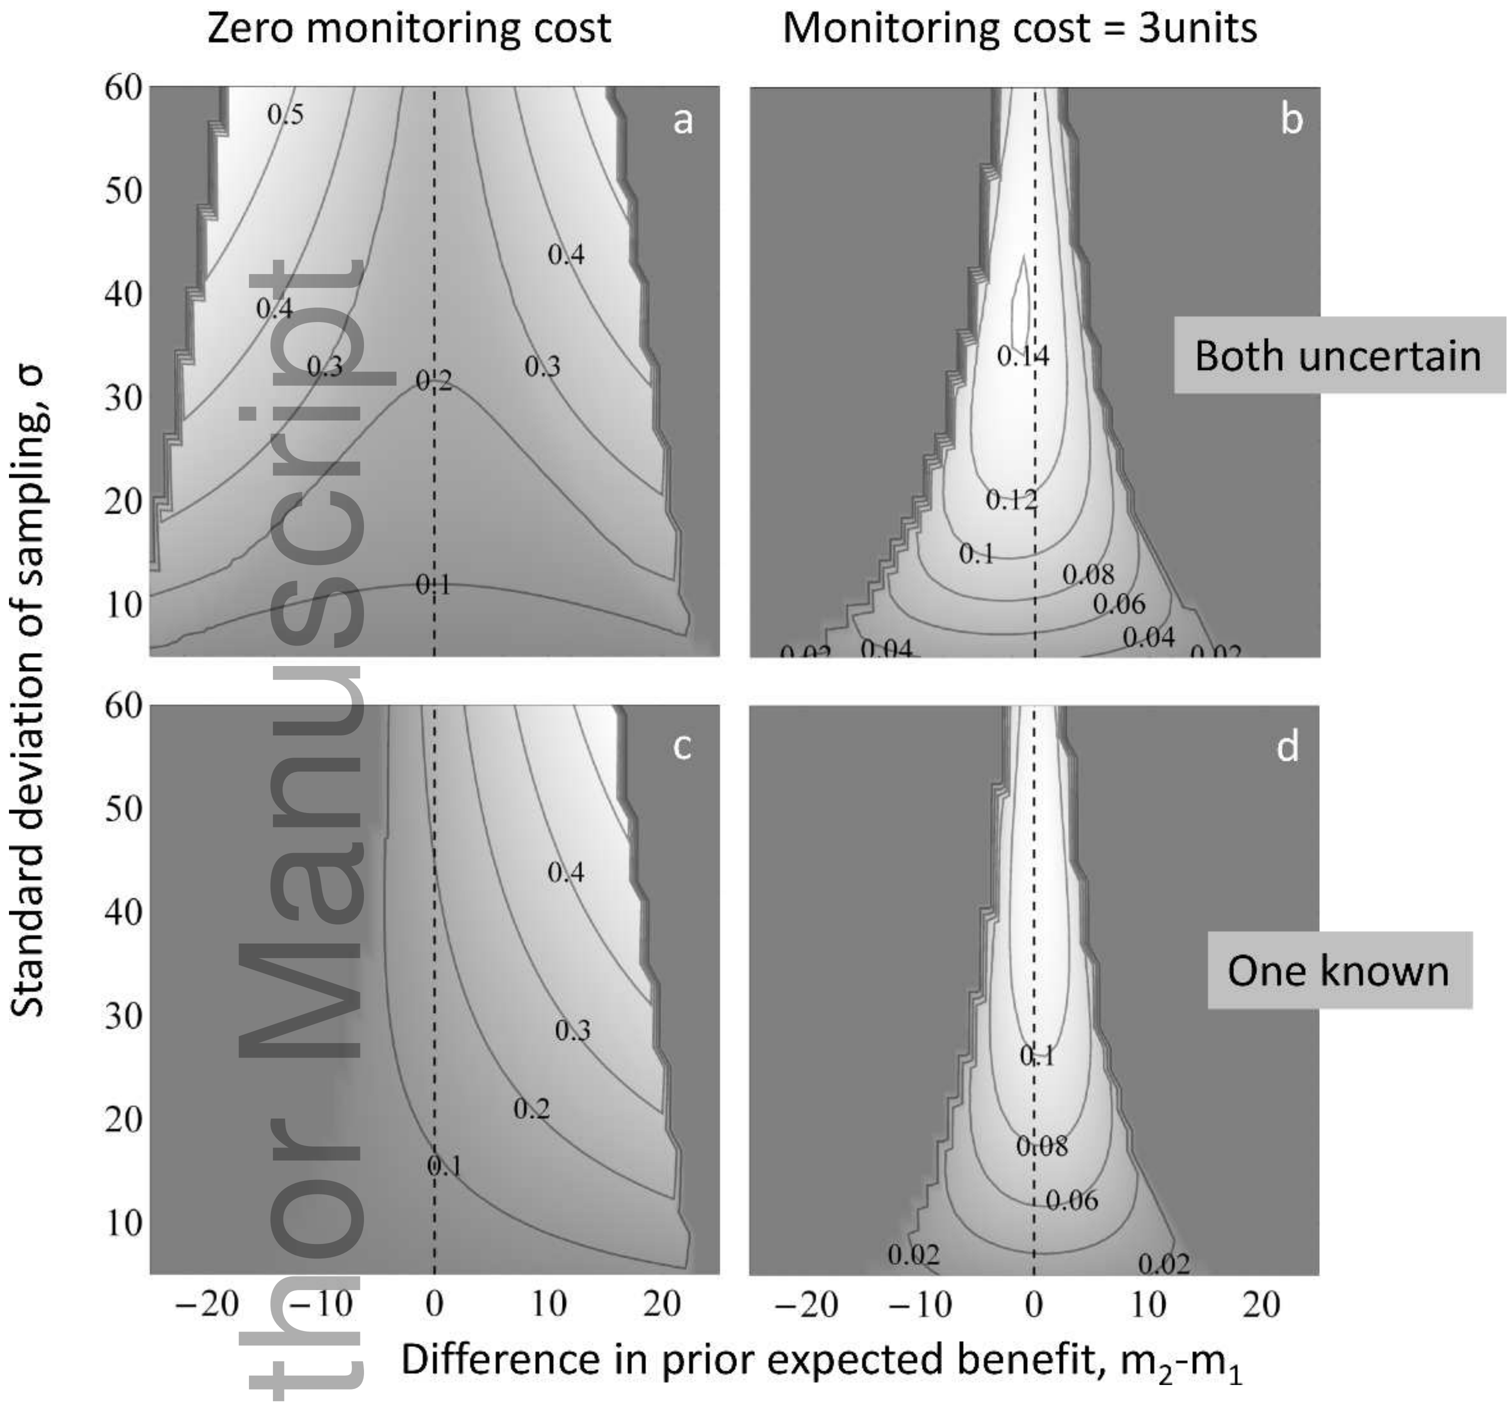

eap_1515_f1.tif

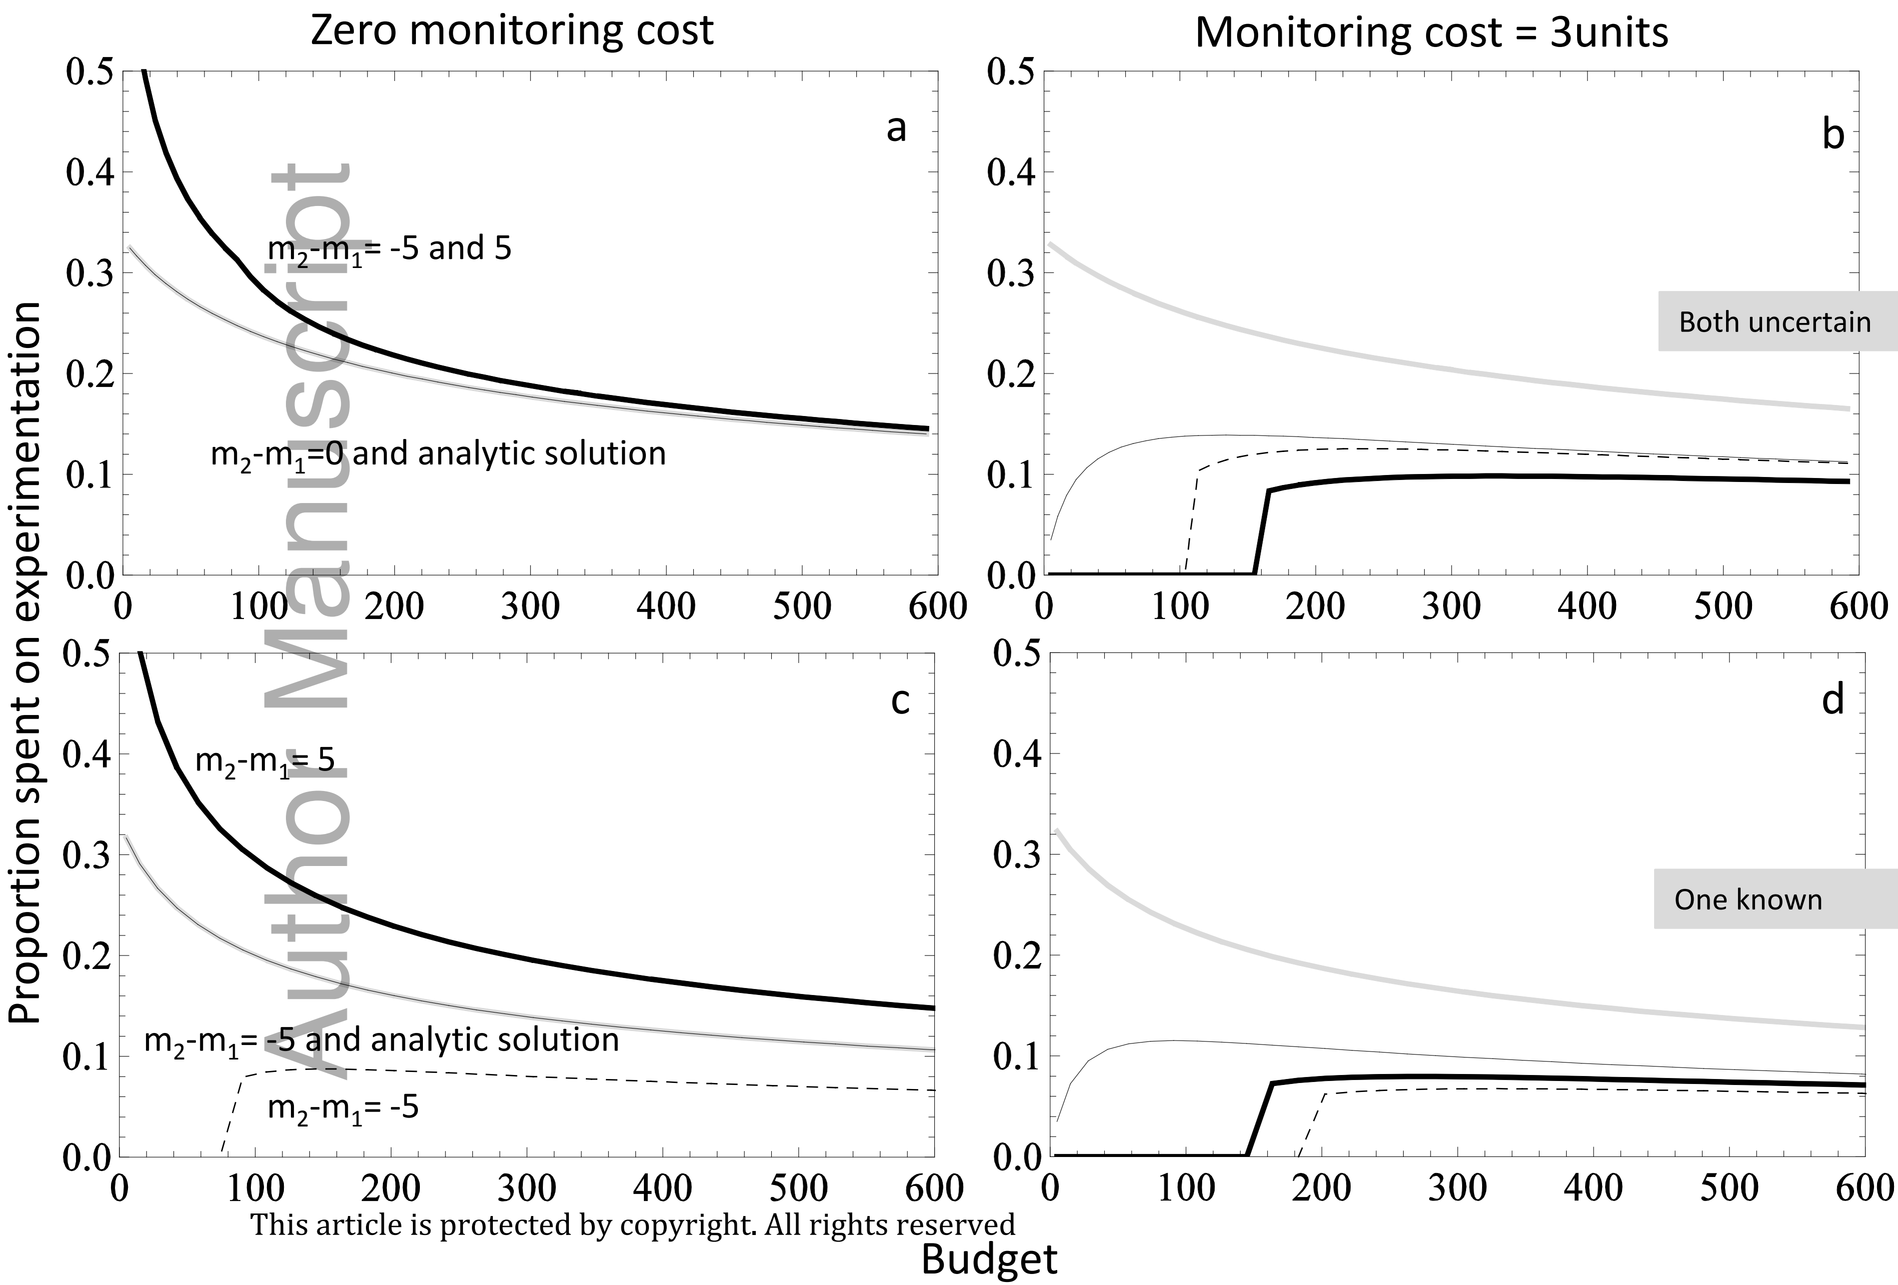
Zero monitoring cost
Monitoring cost = 3units
Proportion spent on experimentation
Budget
a
b
c
d
m2-m1= -5 and 5
m2-m1=0 and analytic solution
Both uncertain
m2-m1= 5
m2-m1= -5 and analytic solution
m2-m1= -5
One known
0.0
0.1
0.2
0.3
0.4
0.5
0
100
200
300
400
500
600

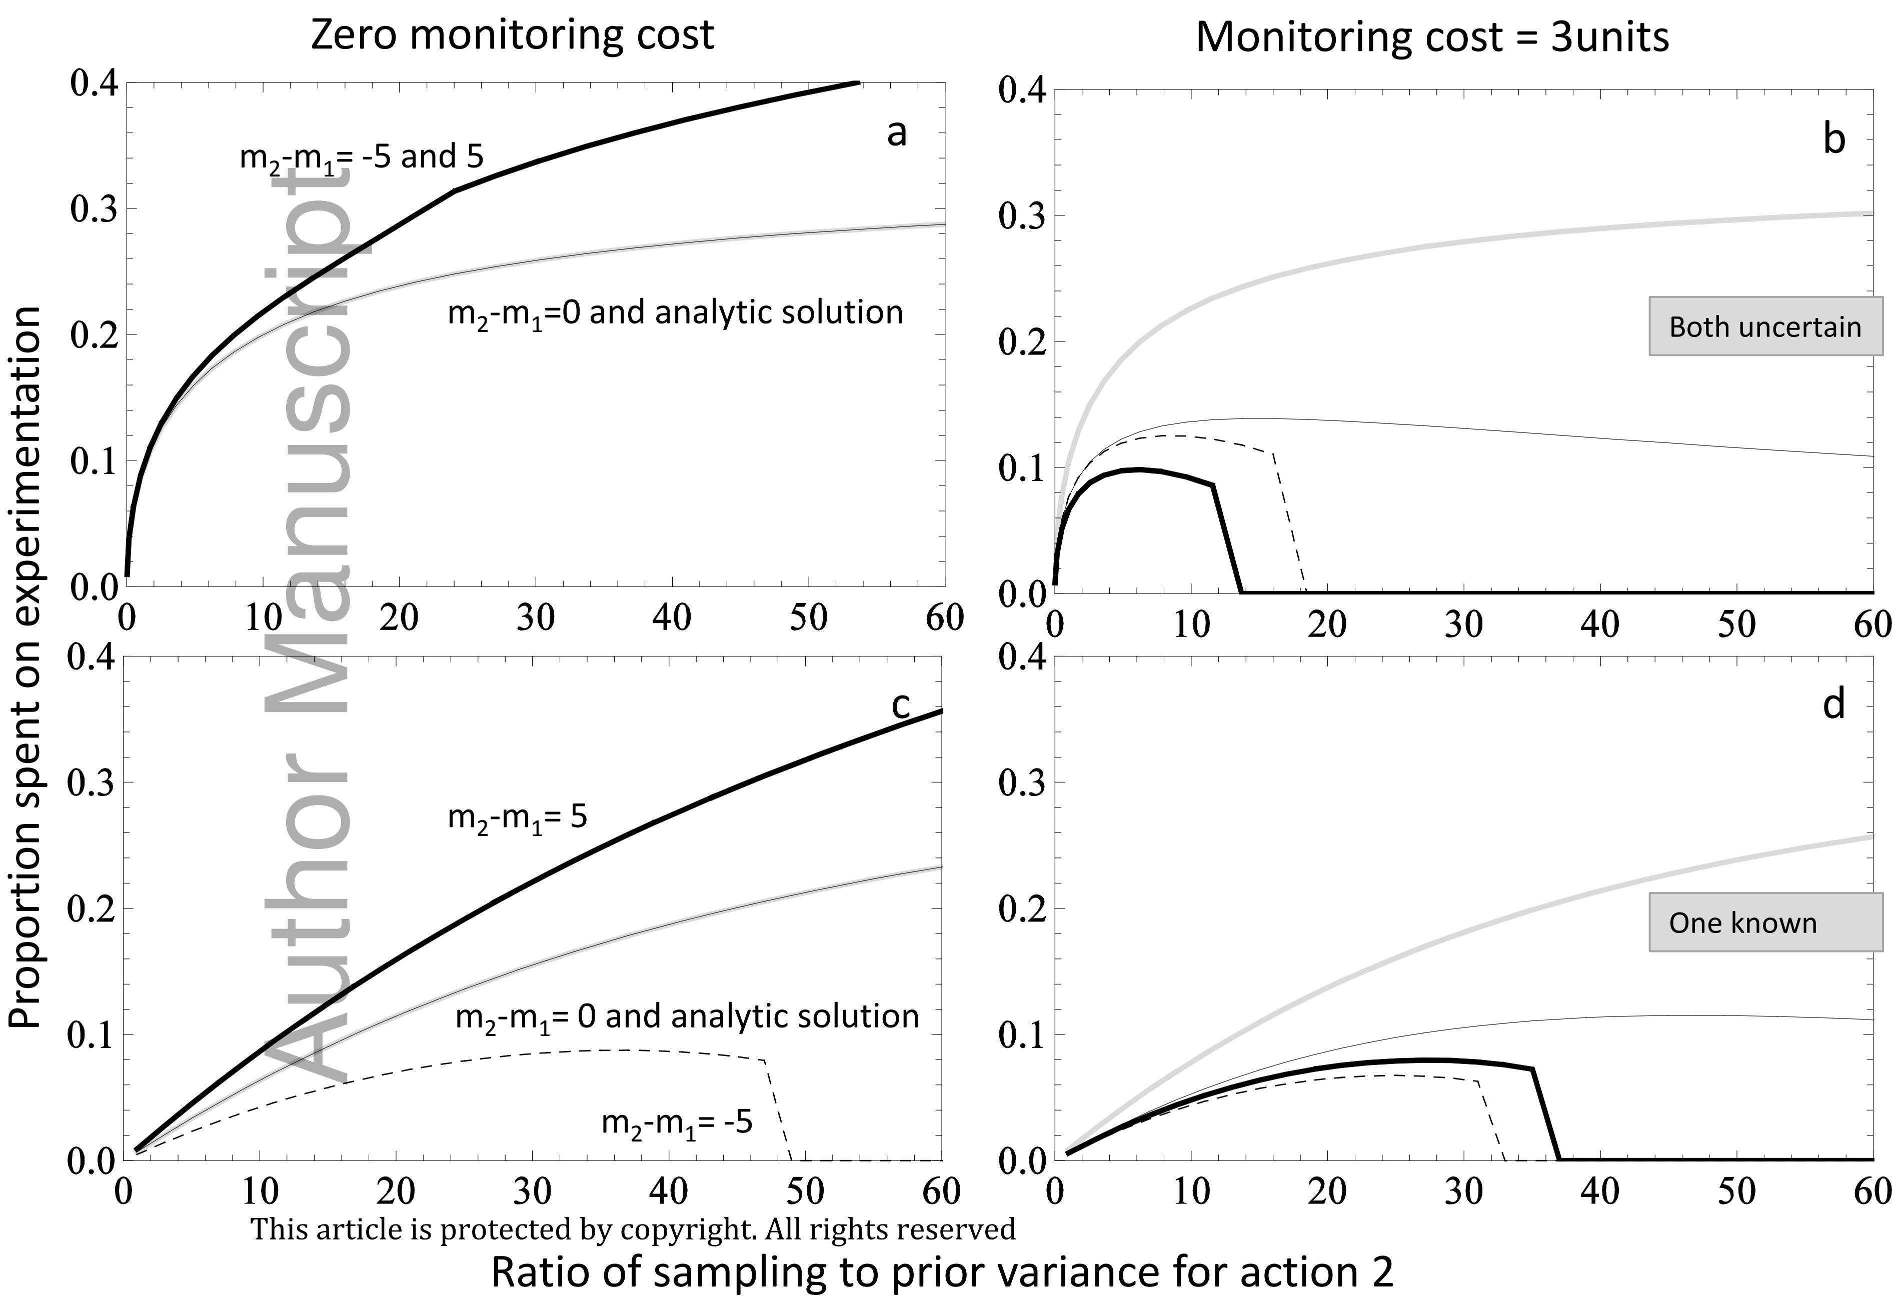
Zero monitoring cost
Monitoring cost = 3units
a
b
c
d
m2-m1= -5 and 5
m2-m1=0 and analytic solution
Both uncertain
m2-m1= 5
m2-m1= 0 and analytic solution
m2-m1= -5
One known
Proportion spent on experimentation
0.0
0.1
0.2
0.3
0.4
0
10
20
30
40
50
60
Ratio of sampling to prior variance for action 2

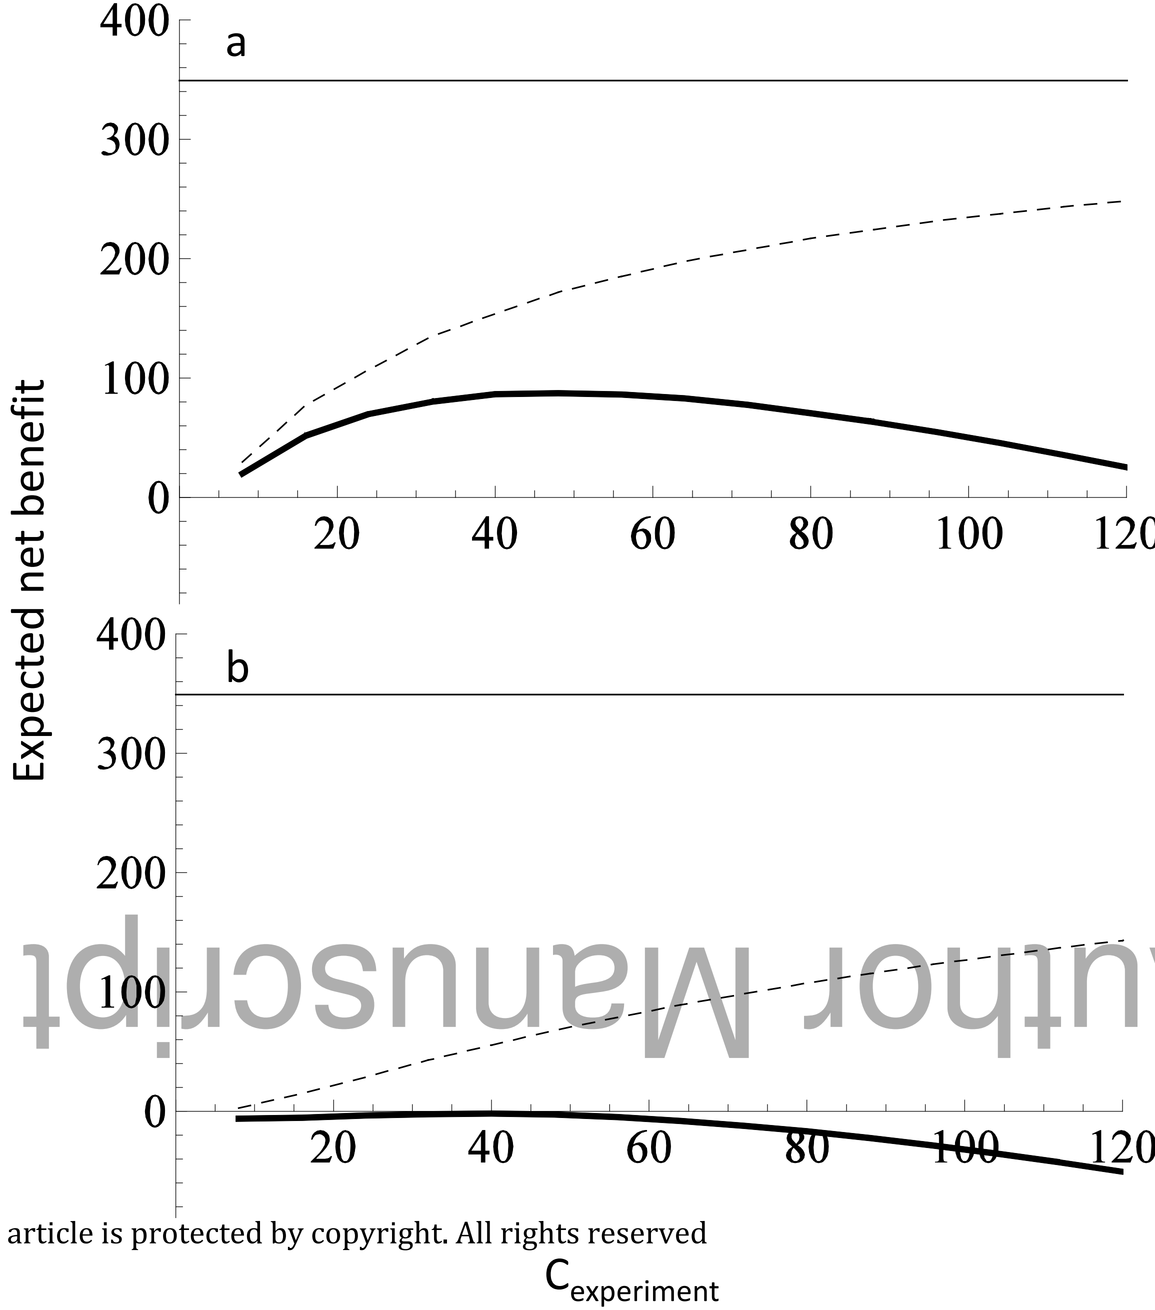
a
b
Expected net benefit
0
100
200
300
400
20
40
60
80
100
120
$C_{experiment}$

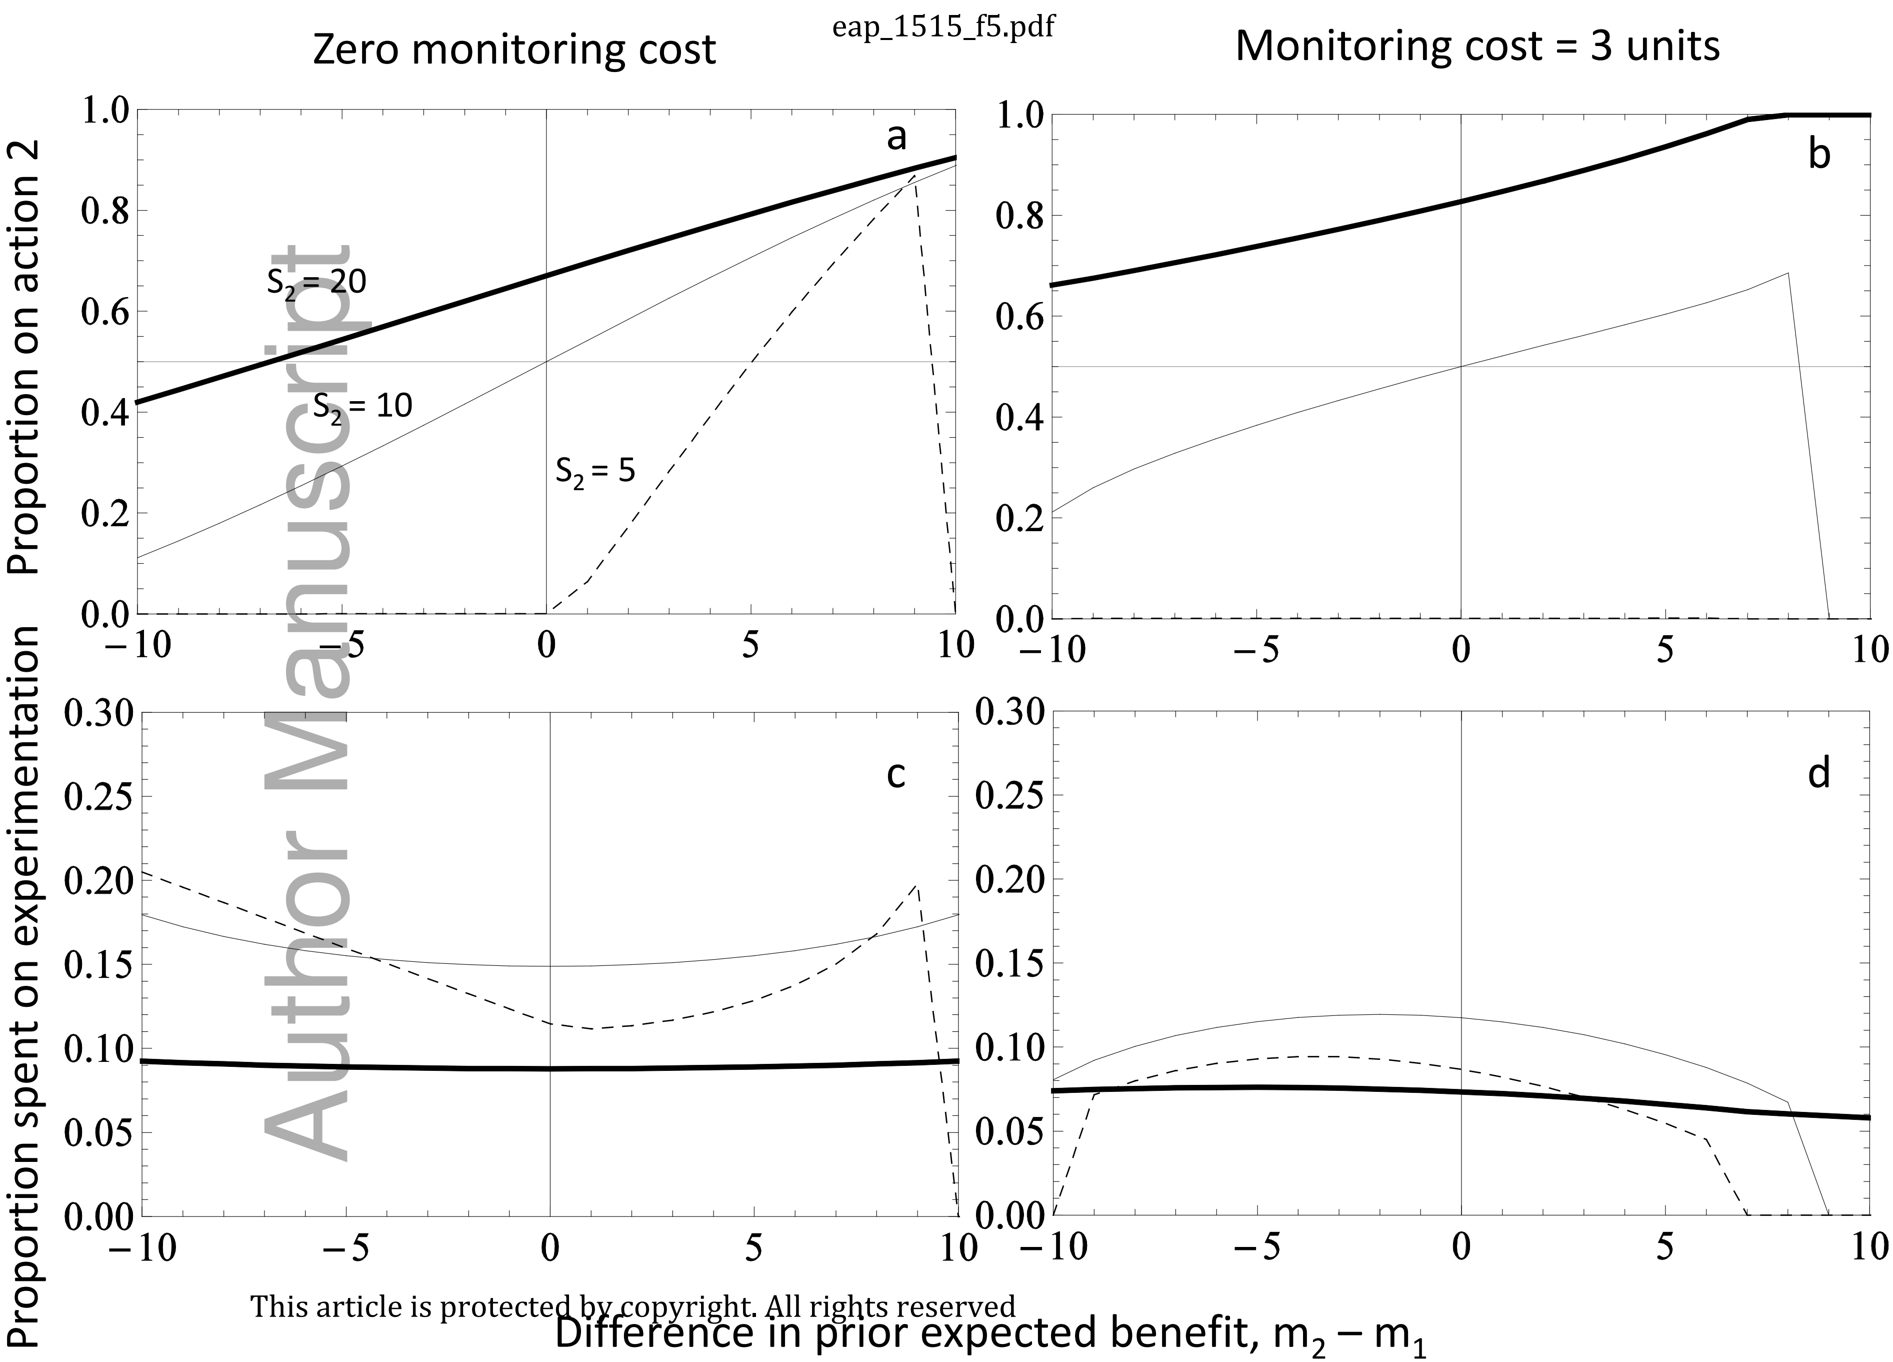
Zero monitoring cost
Monitoring cost = 3 units
Proportion on action 2
Proportion spent on experimentation
a
b
c
d
S2 = 20
S2 = 10
S2 = 5
1.0
0.8
0.6
0.4
0.2
0.0
0.30
0.25
0.20
0.15
0.10
0.05
0.00
−10
−5
0
5
10
Difference in prior expected benefit, m2 – m1

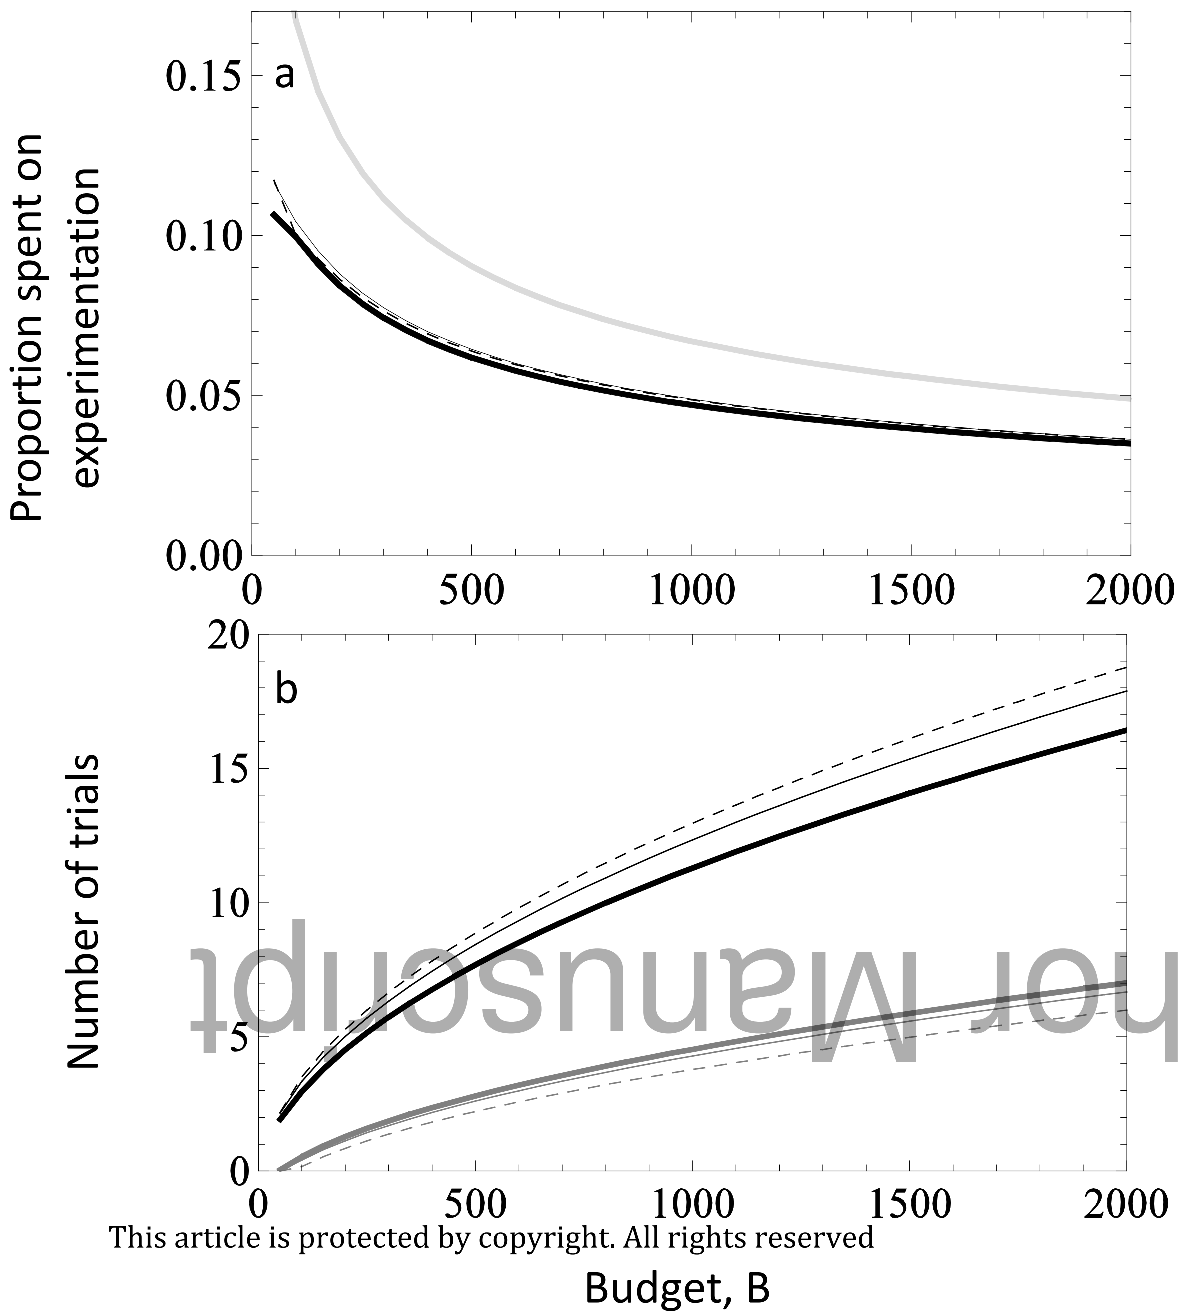
a
Proportion spent on experimentation
0.00
0.05
0.10
0.15
0
500
1000
1500
2000
b
Number of trials
0
5
10
15
20
0
500
1000
1500
2000
Budget, B