Contributed Paper

Improving credibility and transparency of conservation impact evaluations through the partial identification approach

Matthew M. McConnachie,* ¶ Claudia Romero,† Paul J. Ferraro,‡ and Brian W. van Wilgen*

*Department of Botany & Zoology, Centre for Invasion Biology, Stellenbosch University, Private Bag X1, Matieland 7602, South Africa
†Department of Biology, USA & Centre for International Forestry Research, University of Florida, Gainesville, FL, U.S.A.
‡Department of Economics, Andrew Young School of Policy Studies, Georgia State University, Atlanta, GA, U.S.A.

Abstract: *The fundamental challenge of evaluating the impact of conservation interventions is that researchers must estimate the difference between the outcome after an intervention occurred and what the outcome would have been without it (counterfactual). Because the counterfactual is unobservable, researchers must make an untestable assumption that some units (e.g., organisms or sites) that were not exposed to the intervention can be used as a surrogate for the counterfactual (control). The conventional approach is to make a point estimate (i.e., single number along with a confidence interval) of impact, using, for example, regression. Point estimates provide powerful conclusions, but in nonexperimental contexts they depend on strong assumptions about the counterfactual that often lack transparency and credibility. An alternative approach, called partial identification (PI), is to first estimate what the counterfactual bounds would be if the weakest possible assumptions were made. Then, one narrows the bounds by using stronger but credible assumptions based on an understanding of why units were selected for the intervention and how they might respond to it. We applied this approach and compared it with conventional approaches by estimating the impact of a conservation program that removed invasive trees in part of the Cape Floristic Region. Even when we used our largest PI impact estimate, the program's control costs were 1.4 times higher than previously estimated. PI holds promise for applications in conservation science because it encourages researchers to better understand and account for treatment selection biases; can offer insights into the plausibility of conventional point-estimate approaches; could reduce the problem of advocacy in science; might be easier for stakeholders to agree on a bounded estimate than a point estimate where impacts are contentious; and requires only basic arithmetic skills.*

Keywords: adaptive management, evidence based, impact evaluation methods, invasive species, sustainability science

Mejoría de la Credibilidad y la Transparencia de las Evaluaciones de Impacto de la Conservación a través de la Estrategia de Identificación Parcial

Resumen: *El reto fundamental de la evaluación del impacto de las intervenciones de conservación es que los investigadores deben estimar la diferencia entre el resultado después de que una intervención ha ocurrido y entre cuál habría sido el resultado sin la intervención (contrafactual). Ya que lo contrafactual es inobservable, los investigadores deben hacer la suposición incontrastable de que todas las unidades (es decir, organismos o sitios) que no estuvieron expuestas a la intervención pueden ser usadas como sustituto de lo contrafactual (control). La estrategia convencional consiste en hacer un estimado puntual (p. ej.: un número único más un intervalo de confianza) del impacto usando la regresión, por ejemplo. Los estimados puntuales proporcionan conclusiones poderosas, pero en contextos no-experimentales dependen de suposiciones fuertes sobre lo contrafactual, y estas generalmente carecen transparencia y credibilidad. Una estrategia alternativa,*

llamada identificación parcial, consiste en primero estimar cuáles serían los límites contrafactuales si se hicieran las suposiciones más débiles posibles. Después, se reducen los límites al usar suposiciones robustas pero creíbles basadas en el entendimiento de por qué fueron seleccionadas las unidades para la intervención y cómo podrían responder a ella. Aplicamos esta estrategia y la comparamos con las estrategias convencionales al estimar el impacto de un programa de conservación que eliminó a los árboles invasores en una parte de la región florística del Cabo. Incluso cuando usamos nuestro estimado de impacto de identificación parcial más grande, los costos del programa control fueron 1.4 veces mayor que el estimado previamente. La identificación parcial es prometedora para la ciencia de la conservación pues alienta a los investigadores a entender de mejor manera y a tomar en cuenta los sesgos en la selección de tratamientos; puede ofrecer entendimiento de la credibilidad de las estrategias convencionales de estimación puntual; podría reducir el problema de la defensa en la ciencia; podría facilitarle a los accionistas el acordar sobre un estimado limitado en lugar de un estimado puntual en el que los impactos son polémicos; y requiere solamente de habilidades aritméticas básicas.

Palabras Clave: basado en evidencia, ciencia de la sustentabilidad, especies invasoras, manejo adaptativo, métodos de evaluación de impacto

Introduction

Despite growing demands for evaluations of the impacts of conservation interventions, credible evidence of impacts is still scarce (Pattanayak et al. 2010; Miteva et al. 2012). There is limited understanding of how policy actions unfold to influence behavioral change at any point in the policy cycle and how changes on the ground take place (Vedung 1997; Bell et al. 2011). Failure to document what conservation policies work when and why they work has important implications for the dynamics of the conservation policy cycle and ultimately what impact scarce conservation funds have. We describe an approach that addresses the challenge of assessing impacts of conservation interventions in nonexperimental contexts and contributes to ongoing discussions on how impacts should be estimated that inform conservation policy and practice (Baylis et al. 2015).

To measure impact, researchers must compare the outcome after the intervention has occurred with the outcome had it not taken place (the counterfactual) (Ferraro 2009; Angrist & Pischke 2015). For example, the effect of an intervention to control invasive species would be the difference between invasive species presence after the removal treatment (Fig. 1, panel 1c) and what invasive presence would have been if removal had not occurred (Fig. 1, panel 1b). Because the counterfactual is unobservable, researchers have to assume that outcomes from a comparison or control group that was not exposed to the intervention can be used as a surrogate for the counterfactual outcome (e.g., Fig. 2, panel 1b). Thus, the credibility of estimates of conservation impacts weighs heavily on the plausibility of an untestable assumption about the counterfactual.

The most credible and widely used setup to make counterfactual estimates is a randomized experiment. A researcher randomly assigns the intervention to a group of units (e.g., organisms or sites) that receives the treatment and a control group (the counterfactual surrogate) that does not. Given randomization of the treatment, the 2 groups will be, in expectation, similar, with the only difference being that one group is treated and the other is not. There is no bias in treatment assignment (assuming perfect compliance); thus, any difference between the outcomes of the 2 groups can be causally attributed to the treatment or to sampling error. The difference in mean outcomes between the 2 groups provides a point estimate of impact (i.e., a single number along with a confidence interval).

Despite the power of randomization, most conservation programs will have limited ability to use random assignment as part of program implementation. In nonexperimental contexts, researchers have to make far stronger, and hence less credible, assumptions about the counterfactual if they seek a point estimate of the causal effect of the program (Manski & Nagin 1998; Manski 2011); Instead of relying on random assignment, researchers have to select a comparison group based on in-depth knowledge about treatment assignment and systematically account for selection biases (i.e., why some units were treated and others were not [Ferraro & Hanauer 2014]).

An alternative approach, called partial identification (PI), can be used on its own (Hazzah et al. 2014) or in combination with conventional point-estimate methods (e.g., Arriagada et al. 2012; McConnachie et al. 2015). The method sequentially explores the implications of different assumptions on the range of counterfactual outcomes, increasing the strength of the assumptions and assessing how each assumption choice affects the range of impact estimates (the identification region) (Manski 2007). To demonstrate the approach, we assessed the effectiveness and cost-effectiveness of the nonexperimental Working for Water (WfW) program. This program seeks to protect ecosystem services and biodiversity by controlling invasive alien trees in a region of South Africa (van Wilgen et al. 1998).
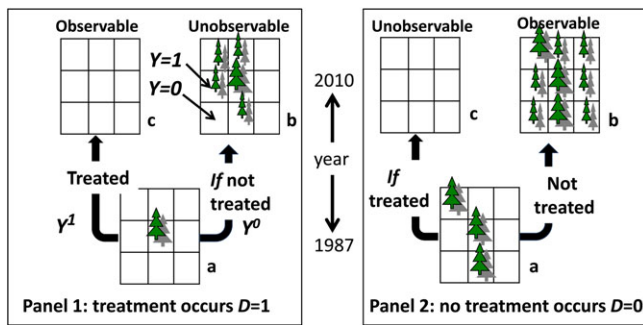
Figure 1. Hypothetical example of the problem of measuring conservation impacts: (a) before and (b) after invasive trees are (c) removed (treated). In panel 1, the invasive tree is removed from the study units and in panel 2 the invasive trees are not removed. Each panel shows 2 potential outcomes: Y^0, not treated (i.e., invasive trees not removed), and Y^1, invasive tree removed. The smaller younger trees in panels 1 and 2 (B) spread from the larger older trees in 1987 (bottom graphs in [a] and [b]).

Methods

Case Study Background

WfW is a national public works program that employs poor South Africans to remove invasive plants (https://www.environment.gov.za/projectsprogrammes/wfw). Established in 1995, the program is arguably the world's most ambitious invasive plant control program (Keonig 2009) and Africa's largest conservation and development initiative with an annual budget of over US$50 million (van Wilgen et al. 1998). Despite the program's magnitude, little is known about its impact (McConnachie et al. 2012, 2013).

We focused on alien tree removal activities within a large (1451 km²) mountainous area of the Cape Floristic Region (CFR) (Fig. 2). The predominant vegetation type is fynbos, a fire-prone shrubland that is susceptible to invasion by fire-adapted invasive trees, even in the absence of human disturbance. Fynbos areas covered by invasive trees consume more water than unoccupied fynbos, resulting in a reduction in water availability (Le Maitre et al. 1996). The CFR contains over 9000 plant species (6210 endemics). Invasion by alien trees is a major risk to the flora, 3087 plant taxa (2972 endemics) are of conservation concern and 1736 taxa (1690 endemics) are in danger of extinction (Raimondo et al. 2009).

In the study area, the long-term goal of WfW is to restore the cover and composition of native fynbos plant species and in so doing to increase water run-off and reduce threats to biodiversity conservation. These goals are achieved by first felling invasive trees and then destroying young regrowth with herbicide spray applications while relying on passive recovery of native vegetation (Holmes et al. 2008). If the young regrowth is not treated before it is approximately chest height, it has to be refelled through costly mechanical clearing methods. Timely and effective follow-up treatments are therefore vital for the long-term success of the program (McConnachie et al. 2012).

Data

We used presence and absence data on invasive trees in cell units of 20 × 20 m for the years 1987 (before WfW) and 2010 (after WfW) for all the units in the study area ($N = 3.70$ million cell units) (Fig. 2 and Table 1). We classified units as treated if they intersected the program's treatment management areas. The management areas comprised about 31.9% of the study area. Treatment records began after 2000 when most of WfW's activities started. Some removal activities, for which we had no records, were carried out by other agencies before when WfW treatments began in the 1990s after the 1987 baseline estimate. We knew that most of these unrecorded activities were located on untransformed protected-area land and state-plantation areas that were abandoned and subsequently cleared in the 1990s (Louw 2004). These were mostly areas that WfW later also treated (see treated–untreated covariate balance statistics in the Supporting Information). Where relevant, for each estimator we discussed how this could result in the treatment effect of WfW being overestimated or the in the case of PI how it could be used to bound impact estimates. Total treatment costs (in South African rands [ZAR]) from 2000 to 2010, including project management costs, were drawn from treatment management data. To account for inflation (approximately 6% per annum), we adjusted all costs to 2014 ZAR. Cost-effectiveness was calculated by dividing the total treatment costs by the reduced area of invasive trees attributable to the WfW intervention.

Defining Conservation Impact

We estimated the difference between observable expected outcomes of treated units and the unobservable expected outcome of these units had they not been treated. In the terminology of impact evaluation, this impact is called "the average treatment effect on the treated" (ATT). The ATT is a common type of impact evaluated in nonexperimental studies (see Ferraro & Hanauer [2014] for other types of impact estimates). In the context of WfW, the expected outcome of the treated group was the expected percentage of units occupied by invasive trees in 2010 for units treated by WfW. The expected counterfactual outcome was the expected percentage of invaded units had these treated units not been treated (Fig. 1).
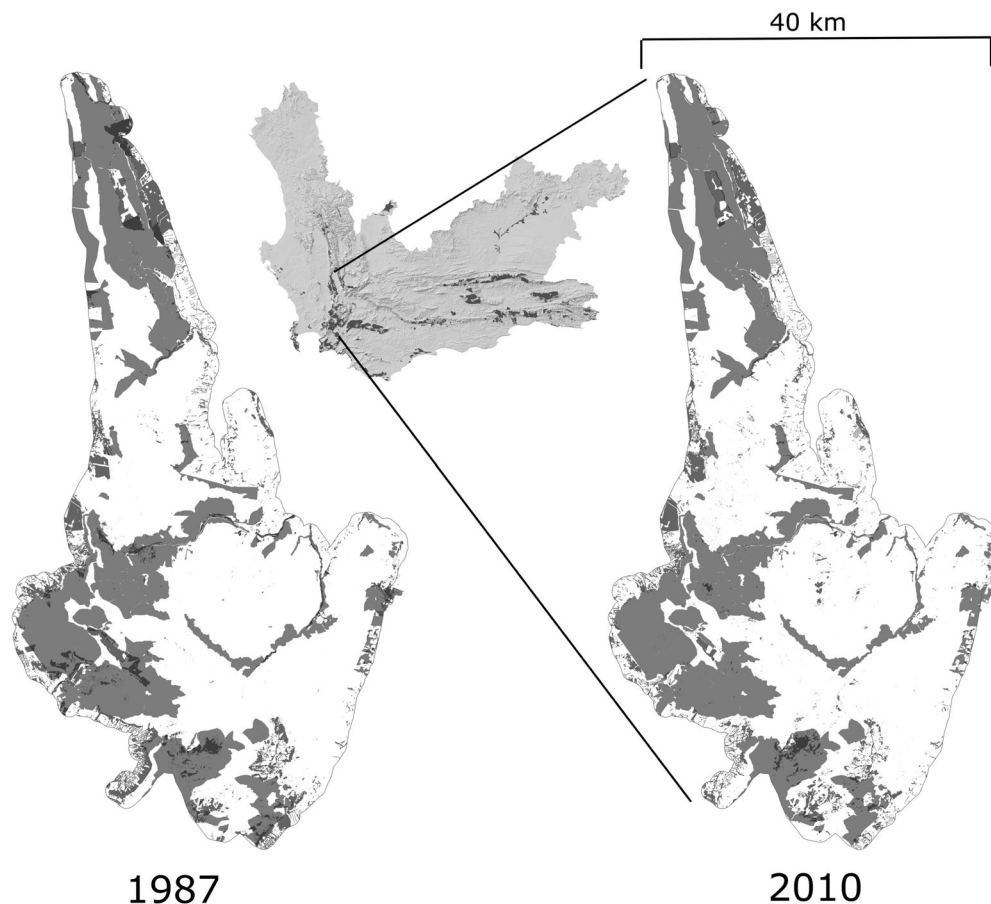
Figure 2. Invasive tree presence marked (dark grey) in 1987 and 2010 in the study area in South Africa's Western Cape Province (inset). Working for Water management areas are marked by large light grey polygons in the main maps and dark grey in the inset.

Table 1. Percentage of study area cell (20 × 20 m) units occupied by invasive trees in treated and untreated units in 1987 and 2010, before and after the implementation of Working for Water in 2000, respectively.[*]

Year	Untreated (SD)	Treated[*] (SD)[*]	Total (SD)
Before 1987	8.75 (28.25)	8.19 (27.42)	8.57 (28.00)
After 2010	8.33 (27.64)	2.87 (16.70)	6.63 (24.88)
Total	8.54 (27.95)	5.53 (22.06)	7.6 (26.44)

[*]*Treatment management areas represented 31.9% of the total study area which covers 1451 km².*

In Fig. 1, panel 1a shows 9 hypothetical units in the study area in the year 1987 that were eventually treated (hereafter denoted $D = 1$) sometime between 2000 and 2010. In 1987, the units had 2 potential outcomes in terms of percentage of units occupied by invasive trees: the potential outcome if the treatment were to happen (panel 1c, hereafter Y^1_{2010}) and the potential outcome were it not to happen (panel 1b, hereafter Y^0_{2010}). Therefore, the ATT in this graphical example

is the expected difference between these 2 potential outcomes for the population of treated units (i.e., panel 1c minus panel 1b):

$$\text{ATT} = E\left[s\left(Y^1_{2010} - Y^0_{2010}\right)|D = 1\right]$$
$$= (0/9 - 5/9) \times 100 = -56\%.$$

Only $(Y^1_{2010}| D = 1)$ (panel 1c) is observable, from which we can estimate $E[(Y^1_{2010}| D = 1)]$. The expected outcome if panel 1a were not treated, $E[(Y^0_{2010}| D = 1)]$, is an unobservable counterfactual (panel 1b). To estimate this expected outcome, one has to make the untestable assumption that the outcome from a comparison group, drawn from units that were not treated (e.g., panel 2b), is a surrogate for the counterfactual outcome of the treated units. Alternatively, a researcher could also model the counterfactual as discussed below (see "Discussion").

PI of Impact and Conventional Approaches

We bounded impact estimates with the PI approach, starting with no assumptions (approach 1) and then adding assumptions about the counterfactual outcome

(monotone treatment selection and response, approaches 2 and 3, respectively). We also considered some conventional point-estimate approaches and their related assumptions: modeling counterfactual outcomes (approach 4), naive treated–untreated comparison (approach 5), before–after treatment comparison (approach 6), before–after-control-intervention comparison (BACI) (approach 7), and conditioning (approach 8). We assessed how credible it would be to make the assumptions of conventional point-estimate approaches given the available data. For each approach, we also determined how the strong assumptions used to make the respective point estimates can be relaxed to make more credible bounded estimates.

A summary of assumptions and estimates of the PI and point-estimate approaches are provided in Table 2 and Fig. 3. We did not include measures of sampling variability (confidence regions) for the PI bounds because we used a population estimate. Confidence regions for PI bounds can be calculated in R and other statistical software (Imbens & Manski 2004).

Results

No Assumptions for PI

As an estimator of $E[(Y^1{}_{2010} \mid D = 1)]$, we used the average percentage of treated units occupied by invasive trees in 2010: 2.87% (Table 1). Without making any assumptions about the unobservable expected counterfactual outcome $E[(Y^0{}_{2010} \mid D = 1)]$, we knew that it could be no higher than 100% (all units would be occupied had WfW not intervened) and no smaller than 0% (no units occupied even if WfW had no intervened). Therefore, with these 2 extreme values, we bounded our ATT estimate within an identification region with the following lower (i.e., best program WfW impact that leads to the greatest possible invasive tree reduction compared with the counterfactual, hence negative value) and upper (i.e., worst program impact compared with the counterfactual) bounds:

$$E\left[\left(Y^1{}_{2010} - Y^0{}_{2010}\right) \mid D = 1\right] \leq \text{ATT}$$
$$\leq E\left[\left(Y^1{}_{2010} - Y^0{}_{2010}\right) \mid D = 1\right]$$
$$(2.87 - 100) \leq \text{ATT} \leq (2.87 - 0)$$
$$-97.13\% \leq \text{ATT} \leq 2.87\%. \quad (1)$$

By having only knowledge of average percentage of treated units occupied by invasive trees, it was possible to narrow the interval width from 200 percentage points without any information (i.e., [−100, 100]) to no greater than 100 percentage points (i.e., [−97.13, 2.87]) (Table 2). The lower-bound (value on the left) implies that WfW caused a 97.13% reduction in invasive tree presence.

The upper-bound implies that WfW caused an increase in invasive tree presence (i.e., implying that, after clearing the units, invasive trees spread faster than they would have if the units had not been cleared) (also see Fig. 3).

Monotone Treatment Selection Assumption for PI

The monotone treatment selection assumption (MTS; Manski & Pepper 2000) is that either positive or negative treatment-selection bias is stronger between the treated and untreated units. Negative selection implies treatment is assigned to units that, in the absence of treatment, have lower levels of invasion, on average, than untreated units. Positive selection implies the opposite relationship between treatment and potential outcomes in the absence of treatment. In some conservation contexts, one can make a credible assumption about the likely direction of selection based on a strong understanding of program implementation; specifically, an understanding of why some units were selected for treatment by the program and others were not (e.g., Arriagada et al. 2012; McConnachie et al. 2015). Factors that cause selection bias (i.e., jointly affect the probability of being treated and the outcome in the treatment's absence) are called confounder variables (Angrist & Pischke 2015).

In our study area, negative selection is plausible. The WfW program in the study region targets untransformed natural areas because of the greater chance of native plant recovery (Holmes et al. 2008) and greater associated expected prevention of native biodiversity loss (Higgins et al. 2001). We know from studies (e.g., Rouget et al. 2003) that higher levels of invasion are associated with human transformed areas that WfW tends to avoid (for empirical patterns consistent with negative selection, see the covariate balance statistics between the characteristics of treated and untreated units in the Supporting Information). The invasive tree presence in 1987 offered support for this assumption: invasive tree presence was lower on treated than untreated units before the treatments happened (Table 1). In addition, we knew that most of the pre-WfW clearing activities were located on WfW treated units. Therefore, the true pretreatment invasive tree presence on treated units is likely to be lower than the 1987 baseline presence.

The assumption of negative selection implies that the expected unobservable counterfactual outcome of treated units is no higher than the observable expected outcome on untreated units, that is $E[(Y^0{}_{2010} \mid D = 1)] \leq E[(Y^0{}_{2010} \mid D = 0)]$. Using observable data from untreated units, we estimated $E[(Y^0{}_{2010} \mid D = 0)]$: 8.33% (based on results in Table 1). Making the plausible assumption of negative selection in the WfW program drastically reduced the lower-bound of the ATT from −97.13 to −5.46:

$$-5.46\% = (2.87 - 8.33)$$
$$\leq \text{ATT} \leq (2.87 - 0) = 2.87\%. \quad (2)$$

Table 2. Summary of approaches used to estimate the impact of the Working for Water project on the reduction of invasive trees.[a]

Approach	Assumption[b]	Assumption illustrated (Fig. 1)[c]	Identification region (estimate bounds)	Advantages	Disadvantages
1. No assumptions	$100 \leq (Y^0_{2010} \mid D = 1) \leq 0$	$100 \leq$ panel 1b ≤ 0	$[-97.13, 2.87]$	only use data, very robust	wide range
2. Monotone treatment selection (− or + selection bias)	$(Y^0_{2010} \mid D = 1) \leq$ or $\geq (Y^0_{2010} \mid D = 0)$	panel 1b $\leq$ or $\geq$ panel 2b	$[-5.46, 2.87]^d$ or $[-97.13, -5.46]$	credible assumptions	range can be wide and includes positive or negative ATT
3. Monotone treatment response	ATT $\geq$ or ≤ 0	panel 1b $\geq$ panel 1c	$[-97.13, 0]^d$ or $[0, 2.87]$	credible assumptions	range can be wide
4. Model counterfactual spread	$(Y^0_{2010} \mid D = 1) = (Y^0_{1987} \mid D = 1) \times 3.75\%$ or 20.6% annual spread	panel 1b = panel 1a × assumed spread rate	−16.3 or −605.6%	precise estimate, allows for temporal changes	could overestimate spread
5. Treated–untreated comparison	$(Y^0_{2010} \mid D = 1) = (Y^0_{2010} \mid D = 0)$	panel 1b = panel 2b	−5.46	precise estimate	seldom plausible
6. Before–after treatment comparison	$(Y^0_{2010} \mid D = 1) = (Y^0_{1987} \mid D = 1)$	panel 1b = panel 1a	−5.32	precise estimate	no possible temporal changes
7. Before–after control intervention	$[(Y^0_{2010} - Y_{1987}) \mid D = 1] = [(Y^0_{2010} - Y_{1987}) \mid D = 0]$	(panel 1b – panel 1a) – (panel 2b – panel 2a)	−4.9%	can control for some time-invariant confounders	cannot adjust for time-variant confounders
8. Conditioning approach	$(Y^0_{2010} \mid D = 1, X) = (Y^0_{2010} \mid D = 0, X)$	panel 1b = panel 2b if confounders balanced (e.g., 1a = 2a)	−2.8%	can control for observable confounders	cannot control for all unobservable confounders

[a]*The first 3 rows use bounded estimates and are therefore called partial identification approaches. The remaining rows present conventional point-estimate approaches that depend on stronger but less credible assumptions.*
[b]*Key: Y^0 and Y^1 are the potential outcome if treated or not, respectively; D = 1 or 0 if actually treated or not, respectively; ATT is the average treatment effect on units that were actually treated, and X represents a set of confounders.*
[c]*The (a) and (b) refer to parts of Fig. 1.*
[d]*Identification region we think is most plausible based on approaches 2 and 3.*
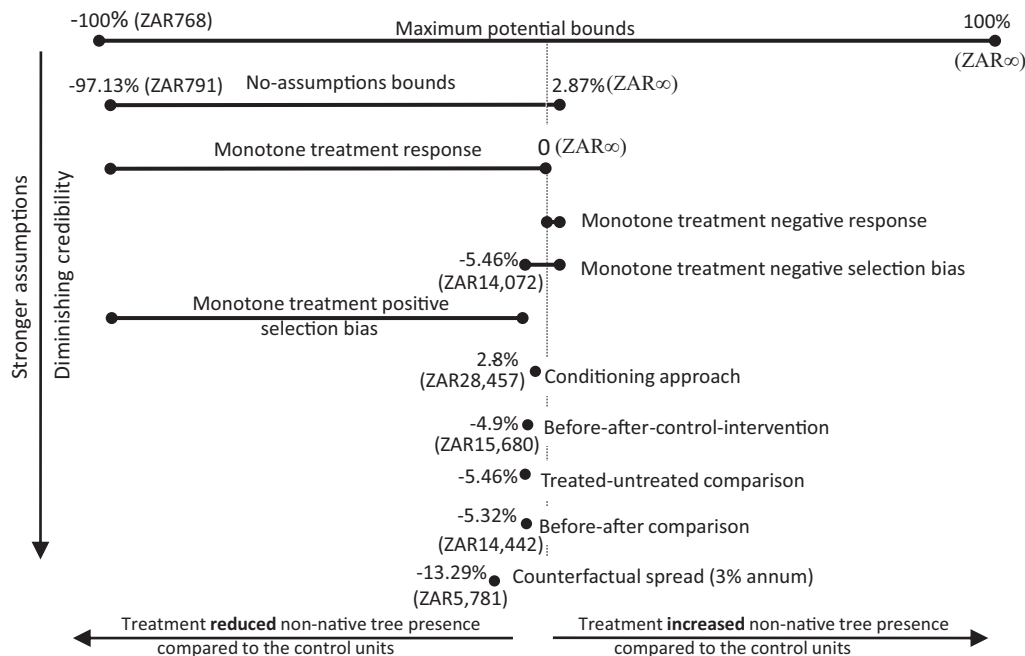

Figure 3. Approaches ordered by the strength of their assumptions used to estimate the impact of efforts by Working for Water to reduce invasive tree presence. Cost-effectiveness of these actions is in parentheses (ZAR per reduced hectare of invasive trees, 10 ZAR = 1 US$, all costs adjusted to 2014 ZAR).

Two sources of positive selection could undermine our MTS assumption. First, WfW usually avoided treating areas far from roads and at high altitudes, which are areas that are less accessible and thus typically less invaded. Second, WfW may have targeted riparian areas because invasive trees consume more water (Holmes et al. 2008). Riparian areas could also be more invaded. But for our MTS assumption to be violated, these 2 potential sources of positive selection would have to be stronger than the well-known sources of negative selection described above. In our study site, the difference between the accessibility and proportion of riparian areas in the treated and untreated groups was not large enough to affect invasive tree presence (Supporting Information); thus, we believe it is plausible to make a negative selection MTS assumption.

Monotone Treatment Response Assumption for PI

The monotone treatment response (MTR) (Manski & Pepper 2000) assumption states that the treatment has either a positive effect on outcomes for all units (impact is zero or positive) or a negative effect. This assumption implies the ATT impact estimate can be greater or less than zero. In our study area, it might be plausible to assume that the effect of WfW on a given unit cannot be positive (i.e., clearing invasive trees would not cause any unit to become occupied by invasive trees). This assumption resulted in the lower bound on the left (best possible program performance) of the ATT shifting from 2.87 to zero. The new range based on only the MTR assumption was

$$-97.13\% = (2.87 - 100) \le \text{ATT}$$
$$\le (2.87 - 2.87) = 0\%. \quad (3)$$

Combining the negative MTS and MTR assumptions yielded the following bounds on the ATT: -5.46% ≤ ATT ≤ 0%.

Our MTR assumption would be implausible, if for example, clearing of nonnative trees had the perverse effect of creating favorable conditions for the establishment of more and potentially new invasive nonnative species that were impossible to control.

Modeling Counterfactual Outcomes for Point Estimation

Some conservation studies model the expected counterfactual outcome based on modeling assumptions about what would have happened in the absence of the treatment from the pretreatment outcome state (e.g., McConnachie et al. 2012). In our case, calculating the counterfactual outcome involved setting the baseline invasion based on data from 1987 (i.e., $Y^0_{1987} | D = 1$) and then multiplying it by an annual spread rate from

1987 till 2010, which would give the counterfactual outcome in 2010 ($Y^0_{2010} | D = 1$). We derive the spread rate from prior research (Higgins, Richardson, & Cowling 2000; Moeller 2010): estimated spread rates between 3.75% and 20.6%. Using these respective spread rates and assuming this counterfactual spread rate holds over 23 years, the ATT equaled

$$\text{ATT}_{3.75\%\text{spread per annum}} = [(2.87) - (8.19'3.75\%^{23})]$$
$$= -16.3\% \text{ and}$$
$$\text{ATT}_{20.6\%\text{spread per annum}} = [(2.87) - (8.19'20.6\%^{23})]$$
$$= -605.6\%. \quad (4)$$

Modeling the counterfactual in this way yielded substantially larger estimates of WfW impacts than were derived using the PI approach and an MTS assumption of negative selection. Assuming higher spread rates made the difference in the 2 estimates even larger. If one believes the rather weak MTS assumption, these modeling estimates are not credible in our study area. The most likely reason invasive trees did not spread as rapidly as the modeled estimate is that other studies assessed spread in the absence of any clearing, whereas in our study area other clearing took place.

Naive Treated–Untreated Comparison for Point Estimation

Unlike the modeling approach, the naïve treated–untreated comparison approach uses data on untreated units as a surrogate for the expected outcome in the absence of treatment for treated units. It assumes there is no treatment selection bias. In other words, the counterfactual expected outcome is assumed to be the same as the expected outcome on untreated units: $E[(Y^0_{2010} | D = 1)] = E(Y^0_{2010} | D = 0)]$. This assumption allows one to make an impact estimate by simply subtracting the average outcome on untreated units from that of treated units.

Using observable data on untreated units (Table 1), we estimated $E(Y^0_{2010} | D = 0)] = 8.33\%$. The resulting point estimate for the ATT was

$$\text{ATT} E\left(Y^1_{2010} - Y^0_{2010} | D = 1\right)$$
$$= (2.87 - 8.33) = -5.46\%. \quad (5)$$

The assumption that $E[(Y^0_{2010} | D = 1)] = E(Y^0_{2010} | D = 0)]$ is very strong and usually impossible to justify in nonexperimental contexts because it implies that treatment assignment is equivalent to random assignment (Ferraro & Hanauer 2014). We expected to see negative treatment selection. Thus, rather than using this approach to generate a point estimate of the ATT, it was more plausible to view it as a lower bound (as derived in approach 3, MTS assumption).

Before–After Treatment Comparison for Point Estimation

This approach assumes that, in the absence of treatment, the expected outcome is the same as it is pretreatment, which allows one to make an impact estimate by simply subtracting the average outcome on treated units before and after treatment. The no-change assumption assumed the counterfactual outcome (e.g., panel 1b, Fig. 1) was the same as before the intervention was implemented (e.g., panel 1a, Fig. 1), that is: $E[(Y^0_{2010}| D = 1)] = E[(Y^0_{1987}| D = 1)]$. Using data from Table 1, $E[(Y^0_{1987}| D = 1)] = 8.19\%$. Therefore, the point estimate of the ATT was

$$\text{ATT } E\left(Y^1_{2010} - Y^0_{1987}\right)|D = 1$$
$$= (2.87 - 8.19) = -5.32\%. \quad (6)$$

In most conservation cases, the no-change assumption is difficult to justify. In our case, we expected the percentage of units occupied by nonnative trees to decline from 1987 to 2010 even in the absence of the WfW program because many nonnative tree plantations in the area were abandoned and clear-cut in the 1990s before our treatment records began. These plantations made up a sizable fraction of the pretreatment, nonnative tree population in the study area.

More sensibly, relative to using the no-change assumption to make a point estimate, we used the no-change assumption to narrow the lower bound of our PI region from -5.46% (based on approach 3, MTS assumption) to -5.32% as calculated above. We therefore made the more credible but weaker assumption that the counterfactual outcome is not greater than the pretreatment outcome on treated units (8.19%).

BACI Comparison for Point Estimation

Also called differences in differences, the BACI approach uses changes over time in the untreated group outcome variable (from pre- to posttreatment) to adjust treatment group changes that would have happened in the absence of treatment; the net change is attributed to the treatment (or sampling error). By assessing changes instead of levels (e.g., approaches 5 and 6), researchers hope to eliminate fixed (time invariant) differences that might otherwise confound impact estimates. For example, a possible confounding factor in our case study was that WfW favors treating state protected land. These areas are less disturbed and invaded than other land ownership types and hence would be less likely to be invaded in the absence of the treatment (Supporting Information). If the effect of this confounder does not change over time, then it can potentially be eliminated by controlling for pretreatment differences in the outcome variable.

The key assumption is that expected change in pre- and posttreatment outcomes for both groups is the same in the absence of treatment (i.e., they move in parallel; called the equal trends assumption): $E[(Y^0_{2010}-Y^0_{1987}) | D = 1] = E[(Y^0_{2010}-Y^0_{1987}) | D = 0]$. To calculate the ATT, average pre and posttreatment expected outcomes are first subtracted from one another for treated and untreated groups. These values indicate respective average changes in the outcome variable for the 2 groups. Second, the resulting group changes are subtracted from one another. For example in Fig. 1, the approach was calculated as follows: (panel 1c - panel 1a) – (panel 2b - panel 2a). In our case study, the ATT was calculated as follows (with values from Table 1):

$$\text{ATT } E\left[(Y^1_{2010} - Y^0_{1987})|D = 1\right]$$
$$-E\left[(Y^0_{2010} - Y^0_{1987})|D = 0\right]$$
$$= (2.87 - 8.19) - (8.33 - 8.75) = -4.9\%. \quad (7)$$

The equal trends assumption depends on 4 strong but often ignored assumptions, some of which also apply to other approaches (Ferraro & Miranda 2014). First, the treatment effect is assumed to be additive and constant (homogenous). This assumption was difficult for us to make because the treated and untreated groups had different levels that likely moderate the treatment effect (Supporting Information). For example, treated areas have a higher proportion of units located in untransformed land, and units with more untransformed land likely have higher treatment effects (i.e., treatment effects are not constant for all units). Second, the functional form is assumed linear. Like the additive treatment effect assumption, this assumption is difficult to make when confounders are imbalanced between treated and untreated units. Furthermore, in most conservation contexts it is difficult to assume that changes over time will be linear. Third, treatment assignment is assumed to be correlated only with fixed characteristics of the land units rather than time-varying characteristics, such as lagged presence of invasive species. Fourth, units respond in the same way to common time shocks. Consequently, we contend the BACI approach is more credible as a lower-bound estimate than a point estimate. Thus, the lower bound (biggest possible reduction) of the ATT increased from -5.46% (with approach 2, MTS PI assumption) to -4.9%.

Conditioning for Point Estimation

Conditioning (e.g., regression or matching methods) is an approach used to account for confounding factors when all confounders can be observed and measured (while approach 7 is usually used when some confounders are unobservable but time invariant). Conditioning works by reweighting the untreated units to select a counterfactual from untreated units that, on average, is balanced in

the distribution of observable confounding factors. This assumption can be formally defined as $(Y^0_{2010} \mid D = 1, X) = (Y^0_{2010} \mid D = 0, X)$, where X is the set of confounding characteristics. Conditioning implicitly assumes that unobservable confounders do not exist or are perfectly correlated with X (Ferraro & Hanauer 2014).

In our case study, observable confounders fell into 5 major categories. The first 2 were related to land ownership and land use (related to preference for working on state protected land that was untransformed, which as mentioned is typically less invaded). The next category included covariates related to the presence and density of invasive trees before treatment. We knew the program tended to treat less dense invasions because they were expected to pose the greatest spread risk (Higgins et al. 2001). The next categories, discussed above, were related to accessibility and where invasive trees grow.

To control for these confounders, we used a matching method in R (Sekhon 2011). With the data in Table 1, we chose a matching algorithm that gave us the best balance: a genetic search optimization algorithm (Diamond & Sekhon 2005) with 1-to-1 matching with replacement. Based on this step, the counterfactual estimate, $(Y_{2010}^{0} \mid D = 1, X)$, equaled 5.5%. The resulting ATT was

$$\text{ATT } E\left[\left(Y^1_{2010} - Y^0_{2010}\right) \mid D = 1\right), X\right]$$
$$= (2.7 - 5.5) = -2.8\% \left(SE\,0.8, p = 0.001\right). \quad (8)$$

If after matching there were still differences in the pretreatment values (e.g., invasive tree presence), the conditioning assumption was less credible. In such situations, the program evaluation literature suggests doing a BACI (approach 7) on the matched sample (Ho et al. 2007). The BACI equal-trends assumption combined with the conditioning assumption would therefore be a weaker and thus more credible assumption than either assumption alone. In our study, there was no difference between the conditioning estimator and the BACI estimator derived from the matched sample because the average pretreatment invasive tree presence value (7.8%) was the same for the treated and matched control units (Supporting Information). Therefore, the ATT had the same value as above: ATT $= E[(Y^1_{2010} - Y^0_{1987}) \mid D = 1]$ - $E[(Y^0_{2010} - ^0_{1987}) \mid D = 1, X] = (2.87 - 7.8) - (5.5 - 7.8) = -2.8\%$ (SE 0.8, $P = 0.001$).

One of the major problems that can undermine the conditioning assumption is that it is difficult to know which and how many confounders to control. With smaller samples only a few confounders can be included, and even for larger samples it can be difficult to find valid counterfactuals (i.e., sufficient balance) for a large number of confounders (King & Zeng 2006).

However, the strong assumptions of the conditioning approach, or any of the other conventional approaches (4)–(8), are not essential. With the PI approach, one can still draw policy-relevant conclusions without having to invoke strong and often less believable assumptions.

Discussion

Deciding What Combination of Assumptions to Maintain

We have shown how the identification of conservation impact does not have to be an all-or-nothing endeavor. Using PI, researchers and practitioners can make valuable inferences even if the comparison group is not a good surrogate for the counterfactual, as is required for conventional point-estimate approaches. In our case study, when we combined the MTR and MTS assumptions (2 and 3), the identification region for the ATT was [−5.46%, 0%] (Fig. 3). The fact that the conditioning estimate (assumption 8) fell within this region strengthens our confidence in the identification region and provides some evidence that the intervention had an impact. Previous studies estimated that it will cost the program at most ZAR10 310 per reduced hectare of invasive trees (Marais & Wannenburgh 2008, adjusted for inflation to 2014-value ZAR). Even when we used the lower bound PI MTS assumption estimate of −5.46%, the resultant cost-effectiveness estimate was still 1.4 times more costly (Fig. 3). Therefore, without having to use strong assumptions, like those used to make the point estimates, we could make a policy-relevant estimate.

Opportunities for Using PI in Conservation Science

We demonstrated how PI, used alone or in combination with conventional point-estimate approaches, offers 6 distinct opportunities for improving the transparency and credibility of nonexperimental conservation impact evaluations. First, it encourages researchers to focus on the most important aspects of nonexperimental impact evaluations: the understanding of why some units are exposed to an intervention and how selection bias can be accounted for. Research in conservation science could benefit from focusing less on statistical modeling methods (e.g., maximizing the precision of estimates) and giving more attention to designs that reduce bias (Armsworth et al. 2009). Second, PI can be used to assess the plausibility of conventional point-estimate approaches and to detect estimates that fall outside the PI bounds established with no assumptions. For instance, both the modeling assumption estimates fell far outside possible bounds. Third, proper use of PI could help reduce the problem of advocacy in science, in which researchers are tempted to select assumptions that fit the conclusions they want (Manski 2011). Presentation to policy makers with a menu of assumptions and attached conclusions (Fig. 1) empowers researchers to better act as honest brokers, although there is always the risk that

decision makers will push actions they already find more attractive (Weiss 1979). Fourth, when conservation impacts are contentious, it might be more amenable for discussions on resource management options to have stakeholders recognize and agree on bounded estimates than on single number point estimates (e.g., Manski & Nagin 1998). Fifth, because PI requires only basic arithmetic skills and knowledge of the intervention context, it could be a valuable tool for researchers and managers who do not have the capacity to use more technically intensive statistical methods such as matching methods. The possibility of incorporating existing knowledge into deciding which assumptions to make about the counterfactual opens doors for transparent and rich interactions among conservation researchers and practitioners.

The design of more effective conservation policies and interventions demands that communities of researchers and practitioners efficiently formalize lessons learned from the past. The definition and framing of what constitutes conservation success is an ongoing process that would benefit from more timely discussions facilitated by approaches such as PI.

Acknowledgments

Supporting Information

Covariate balance statistics (Appendix S1) are available online. The authors are solely responsible for the content and functionality of these materials. Queries (other than absence of the material) should be directed to the corresponding author.

Literature Cited

Angrist JD, Pischke JS. 2015. Mastering metrics: the path from cause to effect. Princeton University Press, Princeton, NJ.

Armsworth PR, Gaston KJ, Hanley ND, Ruffell RJ. 2009. Contrasting approaches to statistical regression in ecology and economics. Journal of Applied Ecology **46:**265–268.

Arriagada RA, Ferraro PJ, Sills EO, Pattanayak SK, Cordero-Sancho S. 2012. Do payments for environmental services reduce deforestation? A farm-level evaluation from Costa Rica. Land Economics **88:**382–399.

Baylis K, Honey-Rosés J, Börner J, Corbera E, Ezzine-de Blass D, Ferraro PJ, Lapeyre R, Persson M, Pfaff A, Wunder S. 2015. Mainstreaming impact evaluation in nature conservation. Conservation Letters XXX:1–7. DOI: 10.1111/conl.12180.

Bell S, Shaw B, Boaz A. 2011. Real-world approaches to assessing the impact of environmental research on policy. Research Evaluation **20:**227–237.

Diamond A, Sekhon J. 2005. Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. Institute of Governmental Studies Working Paper 2006-35, Berkeley, CA.

Ferraro PJ. 2009. Counterfactual thinking and impact evaluation in environmental policy. New Directions in Evaluation **122:**75–84.

Ferraro PJ, Hanauer M. 2014. Advances in measuring the environmental and social impacts of environmental programs. Annual Review of Environment and Resources **39:**397–417.

Ferraro PJ, Miranda JJ. 2014. The performance of non-experimental designs in the evaluation of environmental policy: a design-replication study using a large-scale randomized experiment as a benchmark. Journal of Economic Behavior and Organization **107:**344–365.

Hazzah L, Dolrenry S, Naughton L, Edwards C, Mwebi O, Kearney F, Frank L. 2014. Efficacy of two lion conservation programs in Maasailand, Kenya. Conservation Biology **28:**851–860.

Higgins SI, Richardson DM, Cowling RM. 2000. Using a dynamic landscape model for planning the management of alien plant invasions. Ecological Applications **10:**1833–1848.

Higgins SI, Richardson DM, Cowling RM. 2001. Validation of a spatial simulation model of a spreading invasive plant population. Journal of Applied Ecology **38:**571–584.

Ho D, Imai K, King G, Stuart E. 2007. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. Political Analysis **15:**199–236.

Holmes PM, Esler KJ, Richardson DM, Witkowski ETF. 2008. Guidelines for improved management of riparian zones invaded by invasive plants in South Africa. South African Journal of Botany **74:**538–552.

Imbens GW, Manski CF. 2004. Confidence intervals for partially identified parameters. Econometrica **72:**1845–1857.

King G, Zeng L. 2006. The dangers of extreme counterfactuals. Political Analysis **14:**131–159.

Koenig R. 2009. Unleashing an army to repair alien-ravaged ecosystems. Science **325:**562–563.

Le Maitre DC, van Wilgen BW, Chapman RA, McKelly D. 1996. Invasive plants and water resources in the Western Cape Province, South Africa: modelling the consequences of a lack of management. Journal of Applied Ecology **33:**161–172.

Louw WJA. 2004. General history of the South African forest industry: 1991 to 2002. South African Forestry Journal **201:**65–76.

Manski CF. 1997. Monotone treatment response. Econometrica **65:**1311–1334.

Manski C, Nagin D. 1998. Bounding disagreements about treatment effects: a case study of sentencing and recidivism. Sociological Methodology **28:**99–137.

Manski CF. 2007. Identification for prediction and decision. Harvard University Press, Cambridge.

Manski CF. 2011. Policy analysis with incredible certitude. The Economic Journal **21:**F261–F289.

Manski CF, Pepper J. 2000. Monotone instrumental variables: with an application to the returns to schooling. Econometrica **68:**997–1010.

Marais C, Wannenburgh AM. 2008. Restoration of water resources (natural capital) through the clearing of invasive alien plants from riparian areas in South Africa—costs and water benefits. South African Journal of Botany **74:**526–537.

McConnachie MM, Cowling RM, van Wilgen BW, McConnachie DA. 2012. Evaluating the cost-effectiveness of invasive plant clearing: a case study from South Africa. Biological Conservation **155:**128–135.

McConnachie MM, Cowling RM, Shackleton C, Knight AT. 2013. The challenges of alleviating poverty through ecological restoration: Insights from South Africa's "Working for Water" programme. Restoration Ecology **21:**544–550.

McConnachie MM, van Wilgen BW, Richardson D, Ferraro P, Forsyth A. 2015. Estimating the effect of plantations on pine invasions in protected areas: a case study from South Africa. Journal of Applied Ecology **52:**110–118.

Miteva D, Pattanayak SK, Ferraro PJ. 2012. Evaluation of biodiversity policy instruments: What works and what does not? Oxford Review of Economic Policy **28:**69–92.

Moeller J. 2010. Spatial analysis of pine tree invasion in the Tsitsikamma region, Eastern Cape, South Africa: a pilot study. Honours dissertation, Department of Geography, Rhodes University, Grahamstown.

Pattanayak SK, Wunder S, Ferraro PJ. 2010. Show me the money: Do payments supply ecosystem services in developing countries? Review of Environmental Economics and Policy **4:**254–274.

Raimondo D, von Staden L, Foden W. 2009. Red list of South African plants 2009. Strelitzia **25**. South African National Biodiversity Institute, Pretoria.

Rouget M, Richardson DM, Cowling RM, Lloyd JW, Lombard AT. 2003. Current patterns of habitat transformation and future threats to biodiversity in terrestrial ecosystems of the Cape Floristic Region, South Africa. Biological Conservation **112:**63–85.

Sekhon JS. 2011. Multivariate and propensity score matching software with automated balance optimization: the matching package for R. Journal of Statistical Software **42:**1–52.

van Wilgen BW, Le Maitre DC, Cowling RM. 1998. Ecosystem services, efficiency, sustainability and equity: South Africa's Working for Water program. Trends in Ecology and Evolution **13:**378.

Vedung E. 1997. Public policy and program evaluation. Transaction Publishers, London.

Weiss C. 1979. The many meanings of research utilization. Public Administration Review **4:**426–431.