Peer Review File

Manuscript Title: When and where to protect forests

Reviewer Comments & Author Rebuttals

Reviewer Reports on the Initial Version:

Referees' comments:

Referee #1 (Remarks to the Author):

This paper carries out a global prioritisation analysis of the areas that should be protected to maximise the number of species protected in forested ecoregions. It takes as basic parameters for the optimisation problem the number of species, the amount of existing protected areas, the amount of forest left and the rates of deforestation in each ecoregion. The main result is that one should invest in areas with high return on investment (ROI), which corresponds to the derivate of the species-area relationship (or more precisely, of the species-"protected area" relationship, see Figure S1) divided by the per unit area cost. The paper is well written and in some ways the methods are clearer than the description of the methods in Wilson et al 2006 Nature, on which this paper draws heavily. However I have a major issue with the analysis as I explain below in the first item under general comments. I believe that if that issue is addressed by the authors, this paper could be an important contribution to biodiversity conservation science.

General comments

1. I have doubts about the goal function chosen for optimisation: the sum of the number of protected species at the end of the planning horizon. I believe that a more appropriate goal function would be the number of species at the end of the planning horizon. With the current formulation, it is somewhat trivial that the regions with high current ROI are the ones that should be prioritised base on the simplest heuristic (results as good as 99.4% of the optimum as the paper reports), and one could debate whether the more complex optimisation is useful. It is also relatively trivial that half of the species are protected in the first ten years (Figure 3), as this is simply the response of the species-protected area relationship. The conservation question is arguably how many species will be still present at the end of the planning, so I would use S_iT=(alpha_i*(H_iT-D_iT))^z. Changing the goal function can have impacts on the solution, as areas that have intermediate to high deforestation may then picked up more often, potentially contradicting one of the main results of the paper that threat is not relevant. It may also be possible that such goal function would still produce similar results, but it needs to be demonstrated.

2. The authors seem to assume that all plant species in the forested ecoregion are constrained to forest habitat. This assumption may fail for ecoregions that have a range of forest and non-forest habitats and ignores the possibility of some species adapting to human-dominated habitat (including managed forest). It would be interesting to do a robustness analysis by using something like the countryside species-area relationship.

3. The authors use constant deforestation rates based on current values. Over 50 years the deforestation rates may change, and I am wondering why did the authors not use deforestation rates based on Shared-Socio Economic Pathway scenarios. This of course would not affect much the results for the first 10 years, but may become into play later in the planning horizon.

4. The way in each deforestation in included in the analysis (ie. by shortening the planning horizon) is not well described on lines 73-98. A short explanation here would be helpful.

5. One problem with such analysis is that there is no central global planning authority allocating funding for conservation. Protected areas are mostly set by countries based on national priorities. I am not saying this kind of exercise is not useful, but at least some discussion of how national priority setting may differ (and in reality analysis of this could be done by downscaling the analysis to countries) is warranted. Similarly, GDP based land prices are often not a good proxy for national based allocations, as other factors may come into play in selecting areas for development. The discussion is lacking on literature on these issues.

6. Related to the former, currently the Post-2020 targets are being discussed by the Conv on Biol Div. I wonder if a general % area target per ecoregion (based on % of original natural cover) would be supported by this study or a more individual approach to each ecoregion would be required. I wonder because if the results of the study hold, then regions with little habitat left and little protected area would be the most important to protect to reach such target. Such analysis (even qualitatively) would make this study particularly relevant at this moment.

Specific comments

l.26 Reduction in number of trees is riddled with uncertainty, and a number in the reduction of forest area cover would be more reliable.

l.28 Those 1 millions species extinctions are not being driven by habitat loss alone. A large portion of those would be caused by climate and other drivers.

l.74 Here it is implied that a truly SAR is used, but from the methods it is apparent that in reality only the protected area is being used for the calculation.

6. In l.214 it is stated that if the goal is global conservation a complementarity based analysis would be required. I agree, but it seems that this paper is precisely trying to address global conservation without a complementarity analysis, which is a bit of a contradiction. I wonder if closely located ecoregion may share a large portion of species that would change the results.

Figure 1. I wonder if a map of ecoregions colored by species protection would be better. Right now it is hard to assign circles to regions.

Figure 3. Please use distinct patterns for the different lines

SM
Under alternate species-area exponents make it explicit that U refers to uniform distribution.

Table S5.
I would like to see here stats for ecoregion on amount of area originally protected and amount of area protected after 10 years.

Referee #2 (Remarks to the Author):

The authors present an analysis that identifies priorities for the creation of protected areas to improve plant conservation outcomes. Benefits are quantified using a species-area curve model where area reflects protected forest area, and costs are based on a model of gridded GDP. The framework accounts for rates of deforestation within each ecoregion and existing protected areas.

The primary result highlighted in the abstract (L9) is that prioritising the order of protected area establishment is important for maximising the number of species protected at the end of the planning horizon. Although this conclusion does reflect the analysis, it also highlights one of the fundamental issues: protected area establishment typically occurs within national and subnational jurisdictions, is not stratified by ecoregions, and occurs in parallel among nations - none of which is reflected in this analysis. Hence, it is not clear that this decision support tool informs any real-world problem. The Discussion does not mention any policy mechanisms that this work could inform. As far as I am aware there are no organisations that are prioritising among ecoregions to target PA establishment. Are the authors suggesting that is what needs to occur?

I agree with the perspective that conservation problems are well suited to formulation as scheduling problems. Environmental systems are complex, dynamic, linked biophysical and socioeconomic systems, and accounting for some of those dynamics is likely to make planning more robust. That said, this particular analysis largely ignores those dynamics (e.g. by using parameters that are fixed over time, omitting other key processes, and ignoring uncertainty in state and rate variables), so I do not consider this a strong example of the advantages of formulating a conservation planning as a scheduling problem. In fact, there are so few time-dependent components in this problem formulation that I think it could be solved almost as well using a static formulation. That may be a contrast that the authors could consider adding to better illustrate the value of their scheduling approach.

My other major concerns with this work relate to the assumptions that it makes, which need to be stated and justified more explicitly (noting that several of these assumptions are mentioned in the Discussion):

(1) It assumes that protected area (PA) status is effective in preventing forest loss (the threat of loss impacts only non-protected lands (L364-368)) and promoting species persistence. In many parts of the world this is not the case and a more realistic assessment of the cost of effective protection would need to include costs associated with enforcement (for forest loss prevention) and threat

management (for promoting species persistence).

(2) It assumes that no forest restoration takes place, or that forest restoration provides no benefit to species over the planning horizon. Major forest restoration pledges have been made in the context of agreements like the Bonn Challenge and the NY Declaration (>300 M ha globally). Ignoring that restoration implies this is a worst-case planning scenario, which I would argue is of limited relevance and interest.

There is substantial variation in the proportion of forest remaining among ecoregions, and restoration is a key action in those ecoregions with little remaining forest. Yet there is no mention of restoration in this manuscript. How do the authors suggest readers interpret this analysis in the context of the absence of restoration activities?

(3) The authors maximise "the total number of species within protected forest areas" (L348) but the justification for this is not clear. Doing so implicitly assumes that forested areas outside of protected areas provide no value to species, and that all non-forested areas provide no value to species. These are severe assumptions that require explicit justification. It would be odd if none of the species under consideration had the potential to exist at some density outside of protected forested areas.

(4) It assumes that species richness is a suitable metric to achieve desirable conservation outcomes. The problem with species richness is that it favours the protection of species that are distributed among multiple ecoregions because those species get counted separately for each ecoregion. If a species is well protected in one ecoregion, would that not affect the value of protecting it in another ecoregion. Species richness is a blunt instrument for guiding conservation planning which is why most contemporary conservation planning accounts for species identities explicitly. The authors could argue that species is not the right taxonomic level at which to target conservation and that species within ecoregions is a better option, but they have not done so.

(5) The formulation assumes that the rate parameters are fixed among years. There are some specific problems with this: (a) rates of deforestation are likely to be non-linearly related to the proportion of non-protected forest remaining; (b) land prices are likely to be influenced by the proportion of area protected in an ecoregion, and (c) the degree of uncertainty (variance) in parameter definition is likely to vary substantially among ecoregions implying that accounting for that could have a substantial impact on the solutions identified (partially acknowledged on L187). The sensitivity analyses the authors have conducted are useful for evaluating how the inferences change as the values of the fixed parameters change, but this is not the same thing as formally accounting for uncertainty in the parameters. A stronger approach would have been to specify distributions for those parameters and use robust optimisation procedures.

(6) It assumes that climate change will not impact the distribution of forests over the specified time frame, but climate change is already altering those distributions. The authors need to explain why it is reasonable to ignore climate impacts in a long-term planning exercise.

(7) It assumes that protected area within an ecoregion benefits all species occurring within the ecoregion. Yet ecoregions vary widely in size and, for the larger ecoregions, it is unlikely that all

species would benefit equally from additional protection.

Overall, all models are simplifications of reality, so it is incumbent upon the authors to make it clearer what the simplifying assumptions are and to then justify why the model strikes the right balance between tractability and relevance to the real world. I don't think that has been adequately achieved in this manuscript.

Other comments:

So little information is presented about the species data that it is not possible to evaluate whether it has been used in a suitable manner in this analysis. Providing a source citation for these data does little to justify their use. The Kier et al paper states "For 43% of the 867 ecoregions, data quality was considered good or moderate" (Kier et al 2005), which implies substantial error in richness estimates for many ecoregions. The richness estimates are central to this analysis and the authors need to summarise how those estimates were derived, justify that they are of suitable quality for this analysis, and explain why uncertainty in those estimates can be ignored (I note that none of the sensitivity analyses evaluate the richness estimates).

L176. I think this inference requires qualification. There are a number of NGOs that use benefit:cost analysis (or sometimes cost-effectiveness analysis) to select among potential projects. The conclusion here that ROI is a pragmatic approximation fails to adequately explain that it still requires evaluation of all ecoregions globally in order to achieve that near-optimality. Decision makers could not use ROI on more local or regional levels (as they do now) and still expect to achieve globally optimal results. It is somewhat misleading to describe ROI as having "low data requirements" in this context.

Planning for protection on the basis of plant species richness may not provide solutions that are good for animals. This needs to be acknowledged.

Why is an ecoregion the best spatial unit for this analysis? (The convenience of species data being available at the ecoregion scale is not a sufficient justification). Most decisions about protected areas are made at national scales.

Why was the problem formulated as a budget-constrained problem rather than an ecological-target constrained problem that reported the budget implications? I would argue the latter is more relevant to international policy.

An explanation of what qualifies as a "forested ecoregion" should be added to the SI.

L43-47 These statements do not do justice to contemporary conservation planning. I would argue that for some time best practice in conservation planning has been to account for both opportunity and establishment/operating costs, and for planning solutions to reflect some form of ranking or prioritisation that is not formal scheduling but does at least attempt to provide information that could be used to guide the ordering of priorities in time. Moreover, stating that 'budgets stretch further in less expensive areas' is a poor way of representing the importance of ideas around maximising cost-effectiveness and return on investment. Making these sorts of weakly justified generalisations contribute nothing to establishing the merits of this work.

L73-77. Too little information is provided to understand how the species area curves were parameterised. Is A_i a measure of area or a proportion of remaining habitat? The authors must at least provide a summary of how alpha is parameterised - it is unreasonable to expect readers to dig through the supplement of another paper (ref 21) to figure this out. That reference stated:
"The constant α was calculated via: (the surveyed number of species in an area of interest/(area of interest ^ z))."
and was implemented in the context of counts of parcels, not areas, so it is reasonable to expect the authors to provide some basic explanation of how the parameters were derived. The methods are not repeatable without it. Further to this, given that the parameter z appears to be part of the calculation for alpha, the authors should state more explicitly that alpha was recalculated for each of the sensitivity analyses in which z was varied (see the SM).

L132-140. I thought this was a weak explanation of sensitivity to budget. A diminishing rate of return as a function of increasing budget is so commonplace that it barely warrants mentioning. A more interesting question in the context of a scheduling problem is how much more quickly species are protected as the budget increases.

Minor comments:

Title: This title does not adequately reflect that this analysis addresses both where and when to establish additional protected areas.

L6 "we identify global priorities for conserving plant species". I think this is a poor summary if what this analysis does, as it does not make it clear that this is strictly a PA planning exercise.

L10-16. Isn't the dichotomy presented here artificial? In both cases the algorithm is effectively maximising ROI subject to constraints. Why devote this much space in the abstract to explaining the mechanics of the algorithm?

L68-69, L96 grammar / awkward phrasing

L71. If the authors specify which year of the Kummu et al data they use I cannot find it. Mention it here?

L107 suggest: spending -> allocation or prioritisation (this analysis does not describe spending)

L108 The authors refer to "conservation gains" but I may be worth clarifying that these are actually averted losses.

L113 This is the first mention of "developed" but it is not clear to the reader what qualifies as development and how it fits into this analysis. Development may not be the best choice of word as the authors intend it to mean only deforestation, not the development of non-forested lands. I recommend replacing the word development with a more precise term throughout.

L147 The qualifier 'initial' implies that the conservation benefits are not similar over the longer term?

L206. Species ranges do not appear to be relevant to this analysis as ranges to not appear to have been used in the Kier et al analysis to quantify richness.

L356: "and protected" is inaccurate. The authors appear to mean "and protected forest", because protected non-forested areas are not relevant to their formulation.

L378. I don't understand where the "-1" in the exponents in this equality arise from and some other readers may struggle to understand that too. Perhaps worth explaining this more explicitly in the SI?

Fig 1: There is low contrast between blue and green symbols.

Comments on Supplemental Information:

p2 bullet point, typo "iin", and "we use the species-area
relationship to derive iin each ecoregion" does not make sense. To derive alpha_i in each ecoregion? If so, please provide more explicit detail on how the alphas were derived and what their values are. Why not report those alpha values in Table S5, which could be reduced in font size to fit on the page?

Providing 11 maps is good, but readers would benefit from a matrix of statistics that helps to evaluate the similarities or differences among solutions. It is almost impossible to gauge this objectively by flipping between pages to compare maps.

Referee #3 (Remarks to the Author):

Referee report on “Where to protect first in global forest conservation”
by Ian Luby, Steve J. Miller and Stephen Polasky

Overview

The paper focuses on the issue of how we make decisions regarding conservation. To my mind this is a considerably more important research frontier than that addressed in the bulk of conservation research papers.

That said my first reaction was that this might be a somewhat synthetic comparison of two or more academically concocted approaches. I was therefore relieved to see that the authors contrast these with some of the most commonly adopted approaches to decision making. Adding to this their finding that the use of straightforward and relatively available heuristics provide reasonable approximations to optimised solutions and I feel that this paper is both a strong contribution to the literature and a useful guide for practical policy implementation.

Essentially the paper is concerned not with what to conserve but rather when to conserve it; showing that temporal ordering is of crucial importance in a world of rapidly advancing species loss. This provides something that I always enjoy in a paper – the realisation of what in hindsight looks immediately obvious and clearly important; it's a nice result and an important contribution to the literature.

Specific comments

Line # Comment
39-41 Please support the following statement with a reference:

"Conserved lands are not well chosen to maximize biodiversity conservation but consist in part of land that is low priority for human uses."

70-71 "Land cost was determined using globally gridded GDP from Kummu et al."

I had to go and read the Kummu et al. paper to try and make sense of this approach as my initial reaction was that it seemed a very strong assumption that the GDP of a country, divided by its total land area, should give an unbiased estimation of the per hectare cost of land for conservation in that country. I tried to think of examples where that was clearly true but instead find lots of examples of where this might be quite erroneous, especially for small prosperous countries: Singapore is relatively tiny but has a huge per capita GDP implying a high per hectare GDP; Macao, Luxembourg and Brunei are similar examples. The opposite might apply in very large, low GDP countries.

I admit that, looking at the Kummu et al. paper the maps do give a sensible looking spatial distribution – but the assertion that this validly relates to per hectare conservation costs, especially in rural areas still seems one that deserves more justification. This justification is not forthcoming in the Supplementary Information.

Presumably two countries with similar population and GDP but one of which has an IT dominated economy while the other excels in high value agriculture would be assumed to have similar land costs. That seems unlikely to be the case in practice.

My major concern here is that any error arising from this assumption might not be unbiased. Won't this lead to net benefit measures being biased towards very large countries with low population/GDP per capita?

The assertion that land cost is itself a proxy for conservation cost is also something that needs defence.

I feel this approach needs further discussion and justification, at least in the SI.

74 The species–area relationship (SAR) adopts a classic power function form as per Preston (1962). This transformed linear approach is widely adopted in the literature, but has been the subject of serious critique (see Leitner and Rosenzweig, 1997; Lomolino, 2000) and is fundamentally challenged by non-linear, tipping point models which increasingly have empirical support. My own view is that the power function form is acceptable when we are considering changes which do not threaten species loss. However as that exception is precisely the focus of much contemporary research including this paper then I feel that alternative perspectives should at least be acknowledged. I am not asking that the analysis should be reconsidered; the prevalence of the power function SAR means that it is a reasonable approach and this is not the novel contribution of the paper (and that contribution is very important). But some acknowledgement of this issue would be welcome.

References:

Drake, J. and Griffen, B. (2010) Early warning signals of extinction in deteriorating environments. Nature, 467, 456–459 https://doi.org/10.1038/nature09389

Lomolino, M.V. (2000), Ecology's most general, yet protean pattern: the species-area relationship. Journal of Biogeography, 27: 17-26. https://doi.org/10.1046/j.1365-2699.2000.00377.x

Leitner, W.A. and Rosenzweig, M.L. (1997) Nested Species-Area Curves and Stochastic Sampling: A New Theory, Oikos, Vol. 79, No. 3, pp. 503-512 https://www.jstor.org/stable/3546894

Preston, F. W. (1962) The Canonical Distribution of Commonness and Rarity: Part I, Ecology, Vol. 43, No. 2 pp. 185-215

76 The preceding lines discuss the relationship between species conserved and the area of remaining habitat; the species–area relationship (SAR). A line in the paper that jumped out for me was regarding the curvature of this relationship given by parameter z where the authors write "We assume z = 0.2". Although reference is made to Wilson et al (2006), this still seemed a pretty uncompromising statement, so I was pleased to see in the SI that the authors examine a slightly wider set of values for z ranging between 0.1 and 0.3. Nevertheless, even this felt like it needed further justification and so I took a look through Wilson et al (2006) and found it to be similarly opaque in the main paper. The Supplementary Methods to the latter paper simply states that "We employed a z of 0.2, however our results are insensitive to the value of z.". So again the value of z = 0.2 is asserted rather than empirically based.

Subsequent research has shown that the Wilson et al (2006) assumption of a single SAR for all species can be rather simplistic. Matthews et al (2014) show that generalist and specialist species exhibit significantly different SARs. Given that the proportion of generalist to specialist species is unlikely to be constant globally then assuming a single SAR is at best a simplification which needs to be highlighted. I don't think this is a 'killer problem' but it does mean that readers should be somewhat circumspect about treating the detailed results in the SI as definitive. This however does not take away from the very real contribution which the paper makes in terms of highlighting the temporal ordering problem – that is an important issue. Rather I just feel that the authors need to acknowledge that future research should address these real world complexities.

Reference:

Matthews, T.J., Cottee-Jones, H.E. and Whittaker, R.J. (2014), Habitat fragmentation and the species–area relationship: a focus on total species richness obscures the impact of habitat loss on habitat specialists. Diversity Distrib., 20: 1136-1146. https://doi.org/10.1111/ddi.12227

83-91 [my comment re lines 132-138 supersedes some of this – making clear at 83-91 that you examine sensitivity to budget changes would remove much of the issue here]

The choice of budget might appear at first to be a rather trivial issue; surely whether it is $1Billion or $100Billion is irrelevant to the fundamental findings. However, I doubt this is the case; the bigger the budget, the more can be done now and so the less that the negative impact of sub-optimal timing matters. If we had sufficient budget to address all conservation needs this year then the timing issue effectively disappears. Conversely, the lower the annual budget the greater the potential for temporal problems; and the chosen budget is on the low side and arguably might have been selected to highlight the issue.

Now the simple fact is that the hypothetical case above will remain so – this is simply not going to happen and the temporal effect is real. But I feel that it is worth acknowledging the relationship between budget and temporal sub-optimality.

92-100 I felt that the description of the rules used for dynamic optimisation was a little too brief to aid clear comprehension by the reader. What is the unit used here? Is it the "species-dollar" per hectare value suggested at line 78? i.e. are the authors ranking every single separate hectare of potential conservation land globally and ticking them off sequentially so that the first hectare might be in Sumatra while the second is in Madagascar and so on. Or are they taking this on an ecoregion by ecoregion basis? Or is some other rule being used? It wasn't clear to me and this matters because each approach will lead to different outcomes (and different levels of institutional cost and feasibility).

This just needs clarification.

103-104 This states:

"In the first 10 years, with a $1 billion annual budget, the optimal solution involves selecting habitat conservation in 178 out of 458 forested ecoregions"

This is an interesting statement.

First, it goes some considerable way towards answering my previous point (re lines 92-100) but I am going to leave that comment in because I feel it shouldn't be down to the reader to works these issues out – and also it begs the question; why do it this way? Why look at entire ecoregions sequentially, solving the problem there while doing nothing elsewhere? I am not saying that is the incorrect approach – I just feel that the issue and potential trade-offs should be clarified for readers.

Second, the selection of a 10-year initial period is somewhat arbitrary. Looking at the results reported in Figure S3.1 of the Supplementary Information it seems very likely to me that the initial years of intervention (corresponding to the initial budget allocation) have much greater marginal impact on conservation than subsequent years (indeed this is the major thesis and contribution of the paper). Might it not be nice to make the most of this by reporting for more than just the two periods (up to ten years and beyond ten years) shown here? If not in the main paper then this might be a useful illustration for the Supplementary Information which would really bring home to reader (including, hopefully, policy makers) then need for urgent action now. Such an approach might also overcome the visual problem raised by the fact that, as lines 110-112 state, "most ecoregions which eventually are targeted for conservation over the entire 50-year planning horizon are selected within the first 10 years (178 of 205)." Visually this makes those areas targeted after the first ten-year period difficult to discern at first.

106-108 The statement:

"There is also a conspicuous lack of spending in many of the richest regions of the world including the USA, Europe, Japan and Eastern China."

could easily be misconstrued. It's not that these countries are not or should not be spending on conservation; rather it is that the geographical focus of their spending would be best directed at other areas of the world. Indeed, if they really are motivated by the needs of wild species rather than domestic use values or national political gains then their funding should be directed overseas in this initial period as it will deliver much greater improvements in global biodiversity conservation.

113-114 The statement "In 124 of these 178 ecoregions, all land is either conserved or developed by the end of the first 10 years" puzzled me. The term development has not been mentioned before. Can you clarify your meaning here; is there something I have missed in the description of the analysis?

122-123 I do not understand the term "largely" in the following:

"Spatial patterns of early species protection are driven largely by initial conservation return-on-investment (ROI),"

What else is driving results apart from ROI? How is Threat (in Fig 2) defined? Is it brought in to the analysis and if so then how?

127 - 129 "fast-moving development also limits potential species conservation by reducing the remaining forest area that could be protected."

It is not clear that you capture spatially different speeds of development in your analysis – so are you really referring to the speed of conservation and implicitly assuming that development proceeds at the same rate globally?

132-138 Great to see the budget sensitivity analysis here.

It seems to me that there could be an interesting trade-off analysis here, working out the cost-savings afforded by adopting a temporal optimisation approach, as follows:

• Suppose you consider a world where we ignore temporal optimisation (i.e. the present situation) and we compared one low conservation budget $X1 (delivering conservation x1) with a higher budget $Y1 (delivering conservation y1). This leads to an improvement in conservation of y1 – x1 = z1.

• Now consider a world where the conservation budget is fixed at $X but we can now switch on temporal optimisation. This leads to a new level of conservation x2 where x2 > x1.

• Let's suppose that x2 < y1 (it doesn't actually matter what this relationship is as long as x1 > x2, which you have already proven to be the case)

• Now increase the budget from $X1 to $X2 and, using temporal optimisation until x2 = y1 i.e. the conservation improvement is again z1

• The use of temporal optimisation should mean that $X2 < $Y1

• The difference $Y1 - $X2 should give us the value of information provided by temporal optimisation.

By varying the initial budget amount $X2 you could get a relationship between budget and conservation effectiveness under temporal optimisation which might provide a useful policy message.

Showing policymakers that, by using existing data we can very substantially improve the value for money delivered by conservation spending has to be one of the strongest findings of the paper.

149 – 150 I really like the idea of using heuristics. The approach seems resonant of Wilson et al (2006) who might be worth referencing here as well as at the end of the section (as they are at line 159)

151 – 158 Great results which have a strong policy message.

170 – 175 "The highest priority forest ecoregions for conservation globally lie in Melanesia, Southeast Asia, China, the Anatolian Peninsula, the Amazon, and Central America"

I suggest that you here reinforce the message noted in my comments on lines 106-108 above.

Discussion section This raises a number of caveats (some of which I raise above). However, I feel it could usefully be more positive about the importance of the central message of the paper – that the question of when we should conserve is a vital and under-researched issue which turns out to be of great importance to the challenge of addressing biodiversity loss.

The advantages of countries agreeing to act cooperatively in ways which yield the best returns on available resources is also worth stressing.

Somewhat contradicting this - but perhaps with a more realistic view of feasible policy action - is it worth adding that the issue of temporal optimisation is relevant even at the national level.

Finally, I also felt that the brief discussion of Figure 3 within this section should be extended to reinforce the message that simple heuristics such as net benefit analysis can help realty in prioritising action.

SI Comments included in the above review

Referee #4 (Remarks to the Author):

This is an excellent article that will be a very useful contribution to the ongoing effort to find ways to protect biodiversity worldwide. The conception and statistical modeling are original and will be of general interest. The conclusions are robust and well presented so that others can analyze these data in comparison well. The references are full and adequate, and the writing is clear.
I suggest that they add a short paragraph or even a few sentences about climate change. In my own studies, as of the disappearance of moths and butterflies, it is pretty clear that in the quite near future, climate change will become an even more important driver of species and population loss than land use, which of course has been dominant over thousands of years.

Referee #5 (Remarks to the Author):

General Comments

Luby et al. identify conservation priorities for forested ecoregions globally based on information for each of the 400+ forested ecoregions: plant species richness, estimated cost of conservation, degree of protection, and threat. The advance over other recent work is that they do this together with an optimization that considers the sequence of protection through time.

The biggest gains would be in the first 10 years (half of the total species protected) with more moderate gains in the 50 year time horizon as less efficient options are available.

This is an important paper that begins to address the conservation actions in sequence. I think this complementary perspective is useful in addition to the biodiversity analyses based on spatial priorities and irreplaceability, although I also have my reservations about this explained below. I do appreciate all of the sensitivity analyses and the discussion of the limitations of the study. The focus on plants and forested ecoregions is also a nice addition to the studies focused on vertebrates, and certainly warranted from a conservation perspective.

One of my main reservations is the use of species richness/protection at an ecoregion level. For example, the metric 'species protected in first 10 years'. But we don't actually know which species are already protected, and therefore, don't know which species need to be protected in the future. If we have two ecoregions with the same species richness and the same amount of forest in protection, but one ecoregion has reserves in high quality habitat and the other has reserves in species-poor regions, wouldn't these very different scenarios be effectively the same thing for the purposes here? I realize that changing the z value on the species-area curve is investigated in the appendix, but it doesn't seem to address this problem. This limitation is discussed (Lines 219) and I realize that spatially explicit data is not high enough quality for many plant species to be used systematically, but are there species lists from protected areas ? Could GBIF or similar be used to generate lists inside and outside of protection within each ecoregion? Or perhaps GBIF could be used to generate an estimate of overall sampling effort in each bioregion to be factored into the species-area analysis?

My other main reservation was that the paper overlooked some fundamental aspects of biodiversity conservation- namely the importance of irreplaceability and complementarity. We know from multiple spatial prioritization studies on all kinds of taxa at different scales, that solutions that consider complementarity identify very efficient gains for protecting the pool of biodiversity- much more so than ranking places by species richness. For example, Madagascar is usually very important contribution to the global species pool because of its high levels of endemism, but appears much less important here. Perhaps I've been looking at this map for too long, but I wonder if areas that are naturally sparse in forested vegetation are a bit inflated here (northern boreal zone in Canada, the mallee in South Australia/W. Victoria in Australia). I guess I'm struggling to see how thousands of species could be protected in northern Quebec (with very low endemism and species richness), and how it could (or whether it should) rank more highly than protection of species in Madagascar.

While, I agree that the financial realities are important in conservation decisions, I wonder about the use of a GDP-averaged metric to represent the cost of land. Some ecoregions might have inexpensive land from a global perspective, but that doesn't reflect the costs to the local people relying on those areas for persistence.

It's pointed out that maintenance costs of conservation are not considered, but I also wondered about conservation effectiveness.

Minor Comments

Table S5 would be better as a data file, and it would also be nice to see the code for this paper.

It’s slightly confusing from the writing (e.g. abstract L12-13) whether the sequence of initially targeting ecoregions with little forest remaining+big gains in that remaining forest, then targeting other regions is intentionally done or whether it simply arises from the algorithm as the ‘final return-on-investment’ is always higher in these cases with a truncated trajectory (Fig. S1B). Obviously, there is a space limitation, but given this is a main result, perhaps a couple more sentences of explanation in the main text would help. Figure S1a is quite helpful conceptually, and I would think about adding it to the main text as well.

L149-159 I think this surrogacy analysis needs more explanation. What exactly is the conservation value here called ‘optimal species protection’? If the number of protected species is the metric, then ranking by species richness should provide the optimal solution, and there should be some loss in that if costs/threats are taken into account. Especially given the uncertainty in the cost of conservation, shouldn’t the current results be compared to an iterative ranking of ecoregions that maximizes biodiversity protection irrespective of costs? There should be trade-offs in either direction, how much extra would we have to spend to optimally protect species, and how many species do we lose if we budget.

L 214 – why is it either or ? why not look at endemics separately in the analysis?

I recommend considering some of the other global priorities/endemism papers and citing some of them, and perhaps comparing results in terms of priority areas and gains in efficiency:

Pouzols et at. 2014. Nature 18:383 (they also use ecoregions)

Pollock et al. 2017. Nature 546:141

Hobohm et al. 2019. Perspectives in Ecology and Conservation (endemic-area relationships of plants, which could be used here instead of species-area)

Sandel et al 2019 GEB 29:4

Author Rebuttals to Initial Comments:

Dear Dr. Armstrong:

Thank you for the opportunity to resubmit our revised paper, which is now titled "When and where to protect forests" (originally titled "Where to protect first in global forest conservation"). The referees asked us to extend the analysis to include additional factors in setting conservation priorities and expand our discussion of additional issues not addressed, or partially addressed, in the original submission**. In general, we think the referees have raised excellent points and we have undertaken extensive new analyses and expanded our discussion of other points in the revised version of our paper to address their comments.** Chief among the new analyses is the development of a new probabilistic model that allows us to address complementarity of species occurrences among ecoregions, incorporation of land costs that rise with land scarcity, and incorporation of the potential for reforestation. We have also expanded our discussion of key assumptions and included additional sensitivity analyses to show how results change with different assumptions or with use of different data sources.

We feel that we have addressed all of the major comments from the five referees with additional analyses or discussion with only two exceptions, which are one comment made by referees 1 and 2 about the objective function, and one comment from referee 2 about additional analysis of uncertainty. We carefully considered the comment about the objective function. After further thought and supplementary analyses included in the supplementary materials and referenced in the main text, we choose to stick with our original objective of maximizing the number of protected species and we give a thorough explanation of our reasons in our response to referees 1 and 2. In response to comments from referee 2, we undertook new analyses that allowed for changing parameters through time, including using projections from the Shared Socioeconomic Pathways and rising land costs with increased land scarcity, and expanded our sensitivity analyses. However, we did not opt to go as far as referee 2 suggested in terms of introducing uncertainty or using robust control methods as we do not think doing so would shed important insights and has the potential to obscure rather than illuminate the important main results of the paper.

We have also thoroughly edited the paper to keep the overall word count roughly constant despite the inclusion of the additional analyses, results, and discussion. We feel that the editing to streamline the presentation has also significantly improved the paper.

We include detailed point-by-point responses to each referee comment below. We look forward to hearing from you.

Sincerely,
Ian Luby, Steve Miller, and Stephen Polasky

Original Decision Letter:
Dear Mr Luby

First of all, please allow me to apologise for the delay in giving you this decision. Your manuscript entitled "Where to protect first in global forest conservation" has now been seen by five referees, whose comments are attached below. In the light of their advice we have decided that we cannot offer to publish the present manuscript in Nature.

While the referees find your work of interest, they also raise concerns about the rigour of the approach and the strength of the novel conclusions that can be drawn at this stage. We feel that these reservations are sufficiently important as to preclude publication of the present study in Nature, and we are accordingly closing the file.

Having said this, the referees appreciate the potential contribution of the work, in terms of showing the importance of considering the temporal ordering of conservation action, so should further work allow you to address these concerns we would be happy to look at a revised manuscript (unless, of course, something similar has by then been accepted at Nature or appeared elsewhere). At such a time, you can request that the manuscript file be reopened for resubmission by simply sending an email to that effect (stating manuscript number) to nature@nature.com.

Naturally, we would need you to address all the concerns raised during review. There is much to go at, and many suggestions for additional analyses, including changing the goal of the optimization, using parameters that are dynamic over time, demonstrating the relevance to national priority setting, showing how the findings stand in the context of additional conservation approaches, adding additional sensitivity analyses and better justifying several assumptions, and better demonstrating the value of the temporal approach with additional analysis.

It is not clear to us whether the referee concerns can be addressed in a substantive way. But if you feel they can be then we would be happy to look at a revised manuscript.

I am sorry that we cannot be more positive on this occasion. We hope that you will find our referees' comments helpful.

Yours sincerely

Anna Armstrong
Senior Editor
Nature

RESPONSES TO REFEREE COMMENTS
(Referee comments in black type; our responses in blue type)

REFEREE #1
(Remarks to the Author):

We thank the referee for the excellent set of comments, which pushed us to think harder and go further in the analysis. We think the changes motivated by your comments and those of the other reviewers, have significantly improved the paper. We detail how we responded to your comments below (our responses in blue).

This paper carries out a global prioritisation analysis of the areas that should be protected to maximise the number of species protected in forested ecoregions. It takes as basic parameters for the optimisation problem the number of species, the amount of existing protected areas, the amount of forest left and the rates of deforestation in each ecoregion. The main result is that one should invest in areas with high return on investment (ROI), which corresponds to the derivate of the species-area relationship (or more precisely, of the species-"protected area" relationship, see Figure S1) divided by the per unit area cost. The paper is well written and in some ways the methods are clearer than the description of the methods in Wilson et al 2006 Nature, on which this paper draws heavily. However I have a major issue with the analysis as I explain below in the first item under general comments. I believe that if that issue is addressed by the authors, this paper could be an important contribution to biodiversity conservation science.

RESPONSE: Thank you for the encouraging words, as well as the challenges that you laid out. We have greatly expanded our analysis and revised our discussion of issues in light of the comments by you and the other referees. We have done our best to tackle the challenges you laid out in order to deliver an important contribution in biodiversity conservation science.

General comments

1. I have doubts about the goal function chosen for optimisation: the sum of the number of protected species at the end of the planning horizon. I believe that a more appropriate goal function would be the number of species at the end of the planning horizon. With the current formulation, it is somewhat trivial that the regions with high current ROI are the ones that should be prioritised base on the simplest heuristic (results as good as 99.4% of the optimum as the paper reports), and one could debate whether the more complex optimisation is useful. It is also relatively trivial that half of the species are protected in the first ten years (Figure 3), as this is simply the response of the species-protected area relationship. The conservation question is arguably how many species will be still present at the end of the planning, so I would use S_iT=(alpha_i*(H_iT-D_iT))^z. Changing the goal function can have impacts on the solution, as areas that have intermediate to high deforestation may then picked up more often, potentially contradicting one of the main results of the paper that threat is not relevant. It may also be

possible that such goal function would still produce similar results, but it needs to be demonstrated.

RESPONSE: This is a really important comment and we thought long and hard about it. We agree with you that the number of species that exist far in the future – even well beyond the planning horizon – is the appropriate goal for conservation. After all, what really matters is preventing extinction regardless of whether species occur in protected habitat or unprotected habitat. However, in thinking through the problem carefully, there are several undesirable properties and a serious technical issue with using the number of species extant at the end of the planning horizon as the objective. After careful consideration and analysis, we are convinced that it is better to use the number of protected species as the objective, for reasons that we will explain below

First off, though, we note that there are several cases where the two objectives yield the same solution. When the planning horizon is extended to be long enough (725 years in the base model), all forest area in each ecoregion will be either protected or deforested so that there will be no difference between the number of protected species and the number of species overall at the end of the planning horizon. In the analysis, we could focus on really long planning horizons and then use the number of species at the end of the planning horizon as the objective. However, since this happens multiple centuries in the future, it seemed better to us to analyze a shorter time horizon (even 50 years might be viewed as a really long planning horizon by some). Even for shorter planning horizons, there are some ecoregions for which it is optimal to protect all remaining habitat well before the end of the planning horizon. For example, in our 50-year planning horizon, there are 67 ecoregions in which there is no remaining unprotected forest after 10 years. Here again, there will be no difference between the number of protected species and the number of species overall.

For planning horizons that are not sufficiently long (i.e., less than 725 years), some ecoregions will have remaining unprotected forest at the end of the planning horizon. In this case, there is a difference between the objectives of maximizing species and maximizing protected species. When the objective is to maximize the number of species at the end of the planning horizon, there is no value to protecting forest in ecoregions with remaining unprotected forest because doing so will not change the area of forest remaining, and hence will not change the number of species at the end of the planning horizon. Since the planning horizon is somewhat arbitrary and there is value to conservation efforts even after the planning horizon is reached, not giving any value to additional protection in ecoregions with some amount of unprotected forest doesn't accord with intuition. Indeed, when we consider lengthening the time horizon, additional protection during the initial planning horizon has conservation value because it increases the number of species present at the end of the longer planning horizon. Maximizing the number of *protected* species at the end of the planning horizon captures some portion of the expected future value of conservation beyond the planning horizon ("continuation value" or "salvage value" in dynamic optimization problems). We found that using the objective of maximizing the number of protected species closely matched optimal strategies for longer planning horizons. In the SI, we compared the optimal strategy over the first 10 years and the first 50 years for a 50-

year planning horizon with those of longer planning horizons and found them to be quite similar. We show this in the SI for planning horizons of 100, 200, and 725 years.

In addition to better accounting for the continuation value as the planning horizon is extended, the objective of maximizing protected species is well-behaved (convex, continuous) while the objective of maximizing species is not well-behaved (non-convex) because of the issue of zero marginal value to protection while there is unprotecting forest, which then jumps to positive values when unprotected forest drops to zero. This feature poses more than numerical challenges; it also implies the optimal conservation strategy is not uniquely defined, as there is no 'best' way to prioritize spending of any remaining budget among regions in which the marginal value of protection is zero. This non-uniqueness could easily lead to selection of a protection strategy that is seemingly one of many "optimal" solutions during the planning horizon, but that actually limits conservation beyond the planning horizon (due to ignoring continuation value).

We include a discussion in lines 204-211 that summarizes why we chose to maximize the number of protected species at the end of the planning horizon:

"An alternative objective is to maximize the total number of species in forested areas whether or not they occur in protected forests. For sufficiently long planning horizons, that objective and ours are the same since any remaining forest must be protected or it would have been deforested. When we solve the model over longer time horizons (100 years, 200 years, and 725 years; see SI Section S3), species protection after 50 years is almost identical to what we observe in the 50-year horizon by maximizing species in protected forest areas. This suggests focusing on species in protected forest areas is a reasonable proxy for both objectives."

Regarding your observation on the minor role of deforestation or threat, in our analysis, we found that conservation spending early on was almost exclusively to ecoregions that had high initial conservation ROI. Priority ecoregions, those ecoregions in which all remaining forest was protected before beginning to protect other ecoregions, were characterized by having high initial ROI values, as well as having high ending ROI values even though all of the conservation budget was devoted to protecting all remaining forest in these ecoregions. We were somewhat surprised that these regions were not also characterized by having high threat values. It is important to note that this result is not guaranteed to be the optimal strategy. We provide a counterexample where this is not the case in Figure S1 in the Supplementary Information. But high initial conservation ROI is related to high ending conservation ROI, which is what is crucial for the optimization, while threat value is not. Threat still matters. We find that the optimal solution differs from a heuristic that chooses to allocate conservation budget to ecoregions solely on the basis of which ecoregion has the highest current conservation ROI. However, we find that the deviation between the optimal strategy and this heuristic to be minor. We also discuss the reasoning why we find that high initial conservation in influential in determining conservation priorities, but not threat is not, in the results section in lines 117-126:

“Spatial patterns of early species protection are most strongly associated with high initial conservation ROI, while there is little correlation between the rate of deforestation (threat) and species protection (Fig. 2A). The optimal solution seeks to equalize final conservation ROI, which favors prioritizing ecoregions with high initial conservation ROI. Threat could lead to protection in ecoregions with lower initial conservation ROI if the ending ROI is at least as high as in other areas (see SI Figure S1 for an example). This is more likely to occur in ecoregions with little remaining unprotected forest or with species-area curves with little curvature. However, because this combination rarely occurs in our data, we find only a modest role of threat in our application. SI Section S1 further discusses the limited role of threat.”

2. The authors seem to assume that all plant species in the forested ecoregion are constrained to forest habitat. This assumption may fail for ecoregions that have a range of forest and non-forest habitats and ignores the possibility of some species adapting to human-dominated habitat (including managed forest). It would be interesting to do a robustness analysis by using something like the countryside species-area relationship.

RESPONSE: The referee is correct that we have assumed that plant species are forest obligate species and require forest habitat. Clearly some species are generalists and can survive in non-forest areas or in areas with human modification of the vegetation (e.g., managed forests or plantations). However, we don’t have information to allow us to identify which species can tolerate what types of habitat. We mention species surviving in non-forest habitat in the discussion along with several other limitations of our data and approach that would affect the species-area curve in lines 226-229 in the paper.

Though we don’t have the data to formally analyze this issue, we can think through how including non-forested areas affects the species-area relationship and how that would impact our analysis. It is possible that survival in non-forest areas creates a “species floor” on the species-area curve where some species would still persist even if all forested area was deforested. Additionally, one could add non-forest area to the species-area relationship but this area would come with some sort of reduced habitat quality similar to (though presumably more severe than) our reforestation penalty. With either of these modifications, the conservation ROI will decrease because the species-area relationship is likely to be flattened in some way. In general, reductions in conservation ROI in one ecoregion will favor protection in other areas, though the effects will depend on which of the two methods is used to account for non-forested areas (i.e., where on the species-area relationship flattening is most pronounced).Without better global data, however, we cannot explore this explicitly even though we can suggest how it would affect the results. We include a brief discussion of these conjectures in Section S2 of the SI.

3. The authors use constant deforestation rates based on current values. Over 50 years the deforestation rates may change, and I am wondering why did the authors not use deforestation rates based on Shared-Socio Economic Pathway scenarios. This of course would not affect much the results for the first 10 years, but may become into play later in the planning horizon.

RESPONSE: Following your suggestion, we have conducted an additional analysis using data from the Shared Socioeconomic Pathways (SSPs) (and associated Representative Concentration Pathways). For internal consistency, this analysis uses the SSPs to account not only for changing deforestation rates, but also changing land costs (which follow regional GDP per capita trajectories from the SSPs) and conservation budgets (which follows the global GDP per capita trajectory). We used the SSP3-Baseline scenario, which pairs a pessimistic regional conflict scenario with a similarly pessimistic emissions trajectory. We do so for several reasons, as described in the SI. First and foremost, our analysis is most relevant when threats to forest cover are greatest. Second, in some more optimistic SSP scenarios, net forest cover changes can be positive (reforestation rather than deforestation), and often switch sign from decade to decade. Our analysis is intended to identify optimal budget-constrained protection given *gross* deforestation rates. Because the SSP forest cover changes imply *net* changes, the most pessimistic scenario--with little embedded reforestation--is likely to provide net changes that are closest to gross deforestation rates. The details of our analysis are described in an extension of the "*Alternate forest cover and deforestation rates"* subsection of the Supplementary Information, and the results are presented in Figure S.3.9.

The results of this exercise, depicted in Figure S.3.9, reveal two primary patterns. First, as you conjectured, broad spatial patterns in species conservation during the first 10 years (Panel A) are largely the same as our base scenario. Second, there are clear and sensible regional patterns to differences in final species protection, which are most strongly associated with SSP-derived changes in land cost (Panel B). Because the role of threat is secondary to benefit-cost ratios in the optimal solution (Fig. 2), the differences in deforestation rates are not as crucial (see also Figure S3.7 and Figure S3.8). Instead, the relative increase in GDP per capita--which underlies our land cost estimate--in Asia and Latin America, likely drives reductions in species protection in those regions. Growth in cost in those regions outstrips concurrent increases in global conservation budgets.

While the most of these changes appear in the SI, we also direct readers to this additional analysis in the main text (lines 215-219): "We examined model variants that use deforestation, GDP, and budget projections derived from the Shared Socioeconomic Pathways. We find a shift away from conservation in regions projected to have rapid economic growth leading to higher relative land costs (e.g. Asia; see SI Section S3)."

4. The way in each deforestation in included in the analysis (ie. by shortening the planning horizon) is not well described on lines 73-98. A short explanation here would be helpful.

RESPONSE: Thank you for this comment. It made us realize that we were not as clear or precise in describing the model or the optimization methods in the initial submission as we should have been. We have rewritten the description of the model and optimization approach in the Data and Approach section in the text (see lines 63-83) and the Methods section. In the Methods section we formally describe the basic model and add a new section describing the extended model with complementarity, reforestation, and rising costs. We also go into more depth to describe the numerical optimization methods used.

On your specific question about deforestation, we describe the assumption we use to derive the (constant) deforestation rates by ecoregion in lines 52-54 in the text. We formally write out the evolution of deforestation in lines 326-329 in the Methods section, and explain the end date for potential conservation in lines 331-334: "The end date for potential conservation occurs in an ecoregion when all forested land in that ecoregion is either conserved or deforested ($U_{it} = 0$), or the end of the planning horizon is reached, whichever is soonest."

5. One problem with such analysis is that there is no central global planning authority allocating funding for conservation. Protected areas are mostly set by countries based on national priorities. I am not saying this kind of exercise is not useful, but at least some discussion of how national priority setting may differ (and in reality analysis of this could be done by downscaling the analysis to countries) is warranted. Similarly, GDP based land prices are often not a good proxy for national based allocations, as other factors may come into play in selecting areas for development. The discussion is lacking on literature on these issues.

RESPONSE: We agree with you on both points. On the point about national-level conservation planning we have added the following to the discussion (lines 197-204 in the Discussion section):

"We assumed the objective was to maximize the global number of plant species in protected forested areas at the end of the planning horizon. While our approach is equally applicable at smaller scales (e.g., national protected area planning), we conduct a global prioritization to emphasize potential species protection gains with global cooperation. When conservation planning is done at national scales, it is likely there will be a mismatch between conservation opportunities and budgets leading to unequal conservation ROI across regions, generating lower protection than does coordinated global planning[26]."

We added a reference to Pouzols et al. (2014) (reference 26) who demonstrate the inefficiency of national-level conservation planning relative to global-level conservation planning.

We also note the difficulty of getting globally consistent data on costs. Using the globally gridded GDP was the best proxy we could find. We note several issues with the use of this data. We also expanded the discussion to include other factors that could affect conservation costs (lines 229-235):

"For the costs of protecting forest land, we used gridded GDP as the best consistent globally available proxy. Several prior global studies use agricultural land costs, but this data does not include forest land[33]. Issues with gridded GDP include being an incorrect signal of the value of forest land if the value of economic activity is not primarily related to forests or to land use, significant subsistence or cultural use values not reflected in GDP, and not incorporating maintenance and enforcement costs."

6. Related to the former, currently the Post-2020 targets are being discussed by the Conv on Biol Div. I wonder if a general % area target per ecoregion (based on % of original natural cover) would be supported by this study or a more individual approach to each ecoregion would be required. I wonder because if the results of the study hold, then regions with little habitat left and little protected area would be the most important to protect to reach such target. Such analysis (even qualitatively) would make this study particularly relevant at this moment.

RESPONSE: Thank you for this suggestion. We agree this could provide excellent context for our study. In fact, we added a last sentence to the discussion to make this point explicitly (lines 243-247):

"The upcoming meeting of the Convention on Biological Diversity in China to discuss post 2020 biodiversity goals presents an opportunity for world leaders to increase global cooperation to conserve more of the world's forest biodiversity, and our analysis is well-suited to inform those efforts."

In response to your comment, we have analyzed an additional heuristic protection approach based on a uniform percent protection target across ecoregions. Specifically, in each period, the heuristic assigns the marginal conservation dollar to the ecoregion with the lowest percent of forest area protected, subject to the constraint that land is available in that ecoregion. This heuristic does fairly poorly, achieving just 26% of the overall species protection attainable under the optimal solution, mentioned in the text on lines 140-141, and shown in a revised version of Figure 2 (panel B). Intuitively, a focus on percent protection likely concentrates spending on ecoregions where protection falls on a steep part of the species-area curve. That makes such an approach appear attractive. However, some species area curves are steeper than others, and more importantly, land costs vary substantially across ecoregions. As with other heuristics that do not account for costs, the percent area heuristic thus falls well short of the optimal solution (and benefit-cost heuristic) in terms of species protection. While we fully acknowledge there could be other objectives for forest conservation for which equal percent area protection could be desirable, it falls short as a heuristic if the overall goal is total number of species protected across all ecoregions.

Specific comments

l.26 Reduction in number of trees is riddled with uncertainty, and a number in the reduction of forest area cover would be more reliable.

RESPONSE: We agree that there is large uncertainty and have removed the estimate of the number of trees, and instead focused on forest cover as suggested (lines 27-29).

l.28 Those 1 million species extinctions are not being driven by habitat loss alone. A large portion of those would be caused by climate and other drivers.

RESPONSE: We agree with your point. We have removed the reference to the 1 million species loss because it is driven by multiple causes, not just habitat loss.

l.74 Here it is implied that a truly SAR is used, but from the methods it is apparent that in reality only the protected area is being used for the calculation.

RESPONSE: We have extensively rewritten discussion of the methods in both the main text and the Methods section. We now make clear that we use a true SAR to estimate the parameters α_i for the SAR in the Methods section (lines 319-321). We now distinguish between a classic SAR and our species-protected-area curve in the text in lines 65 – 66: The number of species protected in each ecoregion is a power function of the forest area protected in that ecoregion ("species-protected-area curve") and in the Methods.

6. In l.214 it is stated that if the goal is global conservation a complementarity based analysis would be required. I agree, but it seems that this paper is precisely trying to address global conservation without a complementarity analysis, which is a bit of a contradiction. I wonder if closely located ecoregion may share a large portion of species that would change the results.

RESPONSE: Thank you for raising this point. We have extended our model to account for complementarity: species that inhabit multiple ecoregions. Specifically, we introduce a probabilistic model that maximizes the expected number of unique species protected given protection levels and the sets of ecoregions occupied by different 'profiles' of species. Details of this extension are given in the Methods ("Extensions to the base model"). We show the effect of complementarity on optimal conservation in the results section modified Figure 3 (panels A and D) and discuss this in the text on lines 147-155:

"We extended the basic model to analyze how complementarity, reforestation, and land costs that rise with additional conservation affect optimal conservation (Fig. 3). Accounting for patterns of species overlap means there is greater benefit to conserving one representative ecoregion within a biome more completely while protecting less in nearby ecoregions (Fig. 3A). For example, in the Congo Basin, protected area increases dramatically in the Central Congolian Lowland Forests ecoregion, while protection declines in nearby ecoregions. Ecoregions with many endemic species and no closely related ecoregions (e.g., Hawaii) tend to have higher priority when complementarity is considered."

With regard to your conjecture about nearby ecoregions, our results suggest that this is correct. When overlap is accounted for, protection in ecoregions with similar sets of species (e.g., those within a biome) concentrates into a smaller number of 'representative' ecoregions. Area protected also goes up in ecoregions with a more unique set of species. Both results accord with intuition.

Figure 1. I wonder if a map of ecoregions colored by species protection would be better. Right now it is hard to assign circles to regions.

RESPONSE: Thanks for raising this concern. We evaluated many different designs for Figure 1, including the design you suggested. It is hard to convey information, especially for small ecoregions--many of which are where high levels of species conservation occur--in a global map. It tends to be dominated visually by larger ecoregions and, in our opinion, focuses attention on area patterns rather than the protection patterns we hope to convey. We also examined insets as a way to deal with small areas, but there were too many small ecoregions with high variation in protection scattered across too many locations. In the end, we found the version with the circles was the best at conveying the type of global patterns we wanted to show in ways relatively easy to digest. We acknowledge there are limitations of this visualization choice, and as a result, we also include an exhaustive list of protection by ecoregion in the SI as Table S5. We hope that table can serve a complementary role and clarify any ambiguity.

Figure 3. Please use distinct patterns for the different lines

RESPONSE: What was Figure 3 (now panel B of Figure 2) now uses different line types for the different approaches. We have retained color differences as well in case some readers find those easier to distinguish.

SM
Under alternate species-area exponents make it explicit that U refers to uniform distribution.

RESPONSE: Due to feedback on the randomly sampled species-area exponents, we no longer use this approach, and so your (valid) notation suggestion is no longer relevant. We have replaced the uniform draws with heterogeneous species-area curve exponents from Storch, Keil, & Jetz (2012) and Kier et al. (2005).

Table S5.
I would like to see here stats for ecoregion on amount of area originally protected and amount of area protected after 10 years.

RESPONSE: We agree that this is helpful information. We have redone this table (now Table S3) including information on the basic data for ecoregion (number of species, cost, initial forest area, initial forest protected area, deforestation rate), and information on the optimal solution (period when there is no remaining unprotected forest, conservation spending and additional species protected in the first 10 years, and conservation spending and additional species conserved over the 50 year planning horizon.

REFEREE #2
(Remarks to the Author):

We thank the referee for the extensive set of comments, which pushed us to think harder and go further in the analysis. We have done our best to address all of your comments. We have greatly expanded our analysis and revised our discussion of issues in light of the comments by you and the other referees. We believe that responding to your comments, and those of the other referees, has resulted in significant improvements in the paper. We detail how we responded to your comments below (our responses in blue).

The authors present an analysis that identifies priorities for the creation of protected areas to improve plant conservation outcomes. Benefits are quantified using a species-area curve model where area reflects protected forest area, and costs are based on a model of gridded GDP. The framework accounts for rates of deforestation within each ecoregion and existing protected areas.

The primary result highlighted in the abstract (L9) is that prioritising the order of protected area establishment is important for maximising the number of species protected at the end of the planning horizon. Although this conclusion does reflect the analysis, it also highlights one of the fundamental issues: protected area establishment typically occurs within national and subnational jurisdictions, is not stratified by ecoregions, and occurs in parallel among nations - none of which is reflected in this analysis. Hence, it is not clear that this decision support tool informs any real-world problem. The Discussion does not mention any policy mechanisms that this work could inform. As far as I am aware there are no organisations that are prioritising among ecoregions to target PA establishment. Are the authors suggesting that is what needs to occur?

RESPONSE: We agree that terrestrial protected area establishment and enforcement typically occurs within national and subnational jurisdictions. However, many large conservation organizations do establish priorities for spending at larger scales, including some which operate globally (e.g., WWF, CI). That said, even those organizations that work globally may have constraints on where particular sources of funding can be used, and may focus on a set of specific candidate areas for protection rather than ecoregions so that the question of national-level priority setting is relevant. Despite this, we think that our global focus and prioritization among ecoregions is useful for several reasons.

First, while establishment and enforcement occurs within national boundaries, distributed decision making with local/national objectives and budgets will almost certainly lead to lower levels of global species protection than if efforts were coordinated globally. Our interest in this paper is showing how we could maximize species protection if coordination were undertaken. We've added text to clarify this goal in the discussion (lines 198-204):

"While our approach is equally applicable at smaller scales (e.g., national protected area planning), we conduct a global prioritization to emphasize potential species protection gains with global cooperation. When conservation planning is done at national scales, it is likely there will be a mismatch between conservation opportunities and budgets leading to unequal conservation ROI across regions, generating lower protection than does coordinated global planning[26]."

We added a reference to Pouzols et al. (2014) (reference 26) who demonstrate the inefficiency of national-level conservation planning relative to global-level conservation planning. We think it is helpful to have a clear demonstration of what is possible with more global cooperation, which could act as a spur to global negotiations, such as for the Convention on Biological Diversity COP that will take place in Kunming, China in April-May 2022 (see lines 243-247 in the Discussion).

Second, the methodology we use and the intuition derived from it are equally applicable to a national-scale or sub-national-scale problem in which a planner chooses among domestic protection opportunities.

Third, while establishment of protected areas occurs at scales that are often much smaller than ecoregions, our analysis allows the planner to purchase parcels of land at much smaller increments within an ecoregion. Particular candidate protected areas within each ecoregion could be sorted according to their marginal species-per-dollar, and purchases in our model in a particular ecoregion can be seen to simply pick the best areas (by that marginal species-per-dollar metric) first. In a way, our approach then takes the local prioritization problem to have been solved and asks a broader question about how to allocate conservation dollars regionally. Since it is not uncommon for conservation funding from one country (or citizens of it) to contribute to protected areas in another, we argue there is still a strong need for the broader regional allocation problem to be addressed.

I agree with the perspective that conservation problems are well suited to formulation as scheduling problems. Environmental systems are complex, dynamic, linked biophysical and socioeconomic systems, and accounting for some of those dynamics is likely to make planning more robust. That said, this particular analysis largely ignores those dynamics (e.g. by using parameters that are fixed over time, omitting other key processes, and ignoring uncertainty in state and rate variables), so I do not consider this a strong example of the advantages of formulating a conservation planning as a scheduling problem. In fact, there are so few time-dependent components in this problem formulation that I think it could be solved almost as well using a static formulation. That may be a contrast that the authors could consider adding to better illustrate the value of their scheduling approach.

RESPONSE: Thank you for this suite of suggestions. In response, we have extended our analyses to allow for dynamics of certain model components. First, we now consider the case where costs increase with land scarcity. Increased protection reduces available unprotected forest land, leading to increased land costs (Figure 3, panel C and Methods: Extensions to the

Base Model). Second, we consider evolving deforestation rates, land costs, and budgets by leveraging the Shared Socioeconomic Pathways (and associated Representative Concentration Pathways) projections (See SI: "*Alternate forest cover and deforestation rates"* and Figure S3.9).

We opted not to include uncertainty over state and rate variables. While we do not dispute that uncertainty is present, solving a dynamic optimization problem in the presence of uncertainty presents challenges. It would be relatively straightforward (albeit computationally intensive) to extend the model to account for static uncertainty, which is never reduced or resolved, by maximizing expected protection given joint beliefs over a set of parameters. Indeed, in response to one of your comments and comments by other referees about the need to account for complementarity, we developed a probabilistic model, that could be used to solve for optimal expected conservation benefits. We think, however, that this exercise would add complications without much new insight, and therefore likely obscure rather than illuminate our main messages.

In contrast to static uncertainty, an active learning problem in which experimentation can reveal information about both rates and state variables (in addition to any passive learning that occurs) is the proper approach to making a truly dynamic model with uncertainty about state and rate variables. Such a model would increase the complexity of the problem substantially, requiring a number of additional assumptions (about joint beliefs and learning models, rates of background passive learning, etc.). The curse of dimensionality would also likely render such a problem computationally intractable with 458 ecoregions. So while this is an exceptionally interesting problem, we think this goes well beyond what can be done currently.

My other major concerns with this work relate to the assumptions that it makes, which need to be stated and justified more explicitly (noting that several of these assumptions are mentioned in the Discussion):

(1) It assumes that protected area (PA) status is effective in preventing forest loss (the threat of loss impacts only non-protected lands (L364-368)) and promoting species persistence. In many parts of the world this is not the case and a more realistic assessment of the cost of effective protection would need to include costs associated with enforcement (for forest loss prevention) and threat management (for promoting species persistence).

RESPONSE: We agree that there are additional factors to consider beyond simply establishing protected areas, including management and enforcement costs and management of threats beyond deforestation. If management and enforcement costs were available globally, these recurring costs could be capitalized and included along with land costs to represent the complete costs of conservation. Accounting for these costs is therefore largely a (lack of) data issue. We mention not including maintenance and enforcement costs, along with mention of other cost-related issues, in the Discussion in lines 232-235:

“Issues with gridded GDP include being an incorrect signal of the value of forest land if the value of economic activity is not primarily related to forests or to land use, significant subsistence or cultural use values not reflected in GDP, and not incorporating maintenance and enforcement costs.”

In response to comments about accounting for complementarity, we developed a probabilistic model. A possible extension of our approach could incorporate not only the costs of enforcement and threat management but also the probability of their success.

(2) It assumes that no forest restoration takes place, or that forest restoration provides no benefit to species over the planning horizon. Major forest restoration pledges have been made in the context of agreements like the Bonn Challenge and the NY Declaration (>300 M ha globally). Ignoring that restoration implies this is a worst-case planning scenario, which I would argue is of limited relevance and interest.

There is substantial variation in the proportion of forest remaining among ecoregions, and restoration is a key action in those ecoregions with little remaining forest. Yet there is no mention of restoration in this manuscript. How do the authors suggest readers interpret this analysis in the context of the absence of restoration activities?

RESPONSE: Thank you for this important comment. We agree that restoration is another key tool available to conservation planners. In response to your comment, we have extended our model to allow for restoration of previously deforested land. While there are many ways to model restoration, we consider the costs of restoration to be identical to those for protecting forested land, but assume that restored land contributes less to species richness than conserved primary forest. Details of this extension are in the Methods (“Extensions to the base model”) and its effects on optimal spending are depicted in a modified Figure 3 (panels B and D) and discussed in the text on lines 155-158:

“When reforestation is possible, species gains are further concentrated in priority ecoregions (Fig. 3B) because these ecoregions tend to have higher conservation ROI even though reforested land has lower conservation benefit than primary forest.”

(3) The authors maximise "the total number of species within protected forest areas" (L348) but the justification for this is not clear. Doing so implicitly assumes that forested areas outside of protected areas provide no value to species, and that all non-forested areas provide no value to species. These are severe assumptions that require explicit justification. It would be odd if none of the species under consideration had the potential to exist at some density outside of protected forested areas.

RESPONSE: This is a really important comment, which was also raised by another referee. We thought long and hard about the best way to address it. We agree with you that it would be odd if species in unprotected forests did not contribute to conservation and that it is the number of species that exist far in the future – even well beyond the planning horizon – that is the appropriate goal for conservation. After all, what really matters is preventing extinction regardless of whether species occur in protected habitat or unprotected habitat. However, in thinking through the problem carefully, we found there are several undesirable properties and a serious technical issue with using the number of species extant at the end of the planning horizon in both protected and unprotected forest as the objective. After careful consideration and analysis, we are convinced that it is better to use the number of protected species as the objective, for reasons that we will explain below

First off, though, we note that there are several cases where the two objectives yield the same solution. When the planning horizon is extended to be long enough (725 years in the base model), all forest area in each ecoregion will be either protected or deforested so that there will be no difference between the number of protected species and the number of species overall at the end of the planning horizon. In the analysis, we could focus on really long planning horizons and then use the number of species at the end of the planning horizon as the objective. However, since this happens multiple centuries in the future, it seemed better to us to analyze a shorter time horizon (even 50 years might be viewed as a really long planning horizon by some). Even for shorter planning horizons, there are some ecoregions for which it is optimal to protect all remaining habitat well before the end of the planning horizon. For example, in our 50-year planning horizon, there are 67 ecoregions in which there is no remaining unprotected forest after 10 years. Here again, there will be no difference between the number of protected species and the number of species overall.

For planning horizons that are not sufficiently long (i.e., less than 725 years), some ecoregions will have remaining unprotected forest at the end of the planning horizon. In this case, there is a difference between the objectives of maximizing species and maximizing protected species. When the objective is to maximize the number of species at the end of the planning horizon, there is no value to protecting forest in ecoregions with remaining unprotected forest because doing so will not change the area of forest remaining, and hence will not change the number of species at the end of the planning horizon. Since the planning horizon is somewhat arbitrary and there is value to conservation efforts even after the planning horizon is reached, not giving any value to additional protection in ecoregions with some amount of unprotected forest doesn't accord with intuition. Indeed, when we consider lengthening the time horizon, additional protection during the initial planning horizon has conservation value because it increases the number of species present at the end of the longer planning horizon. Maximizing the number of *protected* species at the end of the planning horizon captures some portion of the expected future value of conservation beyond the planning horizon ("continuation value" or "salvage value" in dynamic optimization problems). We found that using the objective of maximizing the number of protected species closely matched optimal strategies for longer planning horizons. In the SI, we compared the optimal strategy over the first 10 years and the first 50 years for a 50-year planning horizon with those of longer planning horizons and found them to be quite similar. We show this in the SI for planning horizons of 100, 200, and 725 years.

In addition to better accounting for the continuation value as the planning horizon is extended, the objective of maximizing protected species is well-behaved (convex, continuous) while the objective of maximizing species is not well-behaved (non-convex) because of the issue of zero marginal value to protection while there is unprotecting forest, which then jumps to positive values when unprotected forest drops to zero. This feature poses more than numerical challenges; it also implies the optimal conservation strategy is not uniquely defined, as there is no 'best' way to prioritize spending of any remaining budget among regions in which the marginal value of protection is zero. This non-uniqueness could easily lead to selection of a protection strategy that is seemingly one of many "optimal" solutions during the planning horizon, but that actually limits conservation beyond the planning horizon (due to ignoring continuation value).

We include a discussion in lines 204-211 that summarizes why we chose to maximize the number of protected species at the end of the planning horizon:

"An alternative objective is to maximize the total number of species in forested areas whether or not they occur in protected forests. For sufficiently long planning horizons, that objective and ours are the same since any remaining forest must be protected or it would have been deforested. When we solve the model over longer time horizons (100 years, 200 years, and 725 years; see SI Section S3), species protection after 50 years is almost identical to what we observe in the 50-year horizon by maximizing species in protected forest areas. This suggests focusing on species in protected forest areas is a reasonable proxy for both objectives."

(4) It assumes that species richness is a suitable metric to achieve desirable conservation outcomes. The problem with species richness is that it favours the protection of species that are distributed among multiple ecoregions because those species get counted separately for each ecoregion. If a species is well protected in one ecoregion, would that not affect the value of protecting it in another ecoregion. Species richness is a blunt instrument for guiding conservation planning which is why most contemporary conservation planning accounts for species identities explicitly. The authors could argue that species is not the right taxonomic level at which to target conservation and that species within ecoregions is a better option, but they have not done so.

RESPONSE: Thank you for this important comment. In response to this comment and similar comments from other referees we have extended our model to account for complementarity: species may inhabit multiple ecoregions such that sets of species in different ecoregions overlap. Specifically, we introduce a probabilistic model that maximizes the expected number of unique species protected given protection levels and the sets of ecoregions occupied by different 'profiles' of species. Details of this extension are in the Methods ("Extensions to the base model") and its effects on optimal spending are depicted in a modified Figure 3 (panels A and D) and discussed in the text on lines 149-155.

"Accounting for patterns of species overlap means there is greater benefit to conserving one representative ecoregion within a biome more completely while protecting less in nearby ecoregions (Fig. 3A). For example, in the Congo Basin, protected area increases dramatically in the Central Congolian Lowland Forests ecoregion, while protection declines in nearby ecoregions. Ecoregions with many endemic species and no closely related ecoregions (e.g., Hawaii) tend to have higher priority when complementarity is considered."

In general, we find that when overlap is accounted for, protection in ecoregions with similar sets of species (e.g., those within a biome) concentrates into a smaller number of 'representative' ecoregions. Area protected also goes up in ecoregions with a more unique set of species. Both results accord with intuition.

(5) The formulation assumes that the rate parameters are fixed among years. There are some specific problems with this: (a) rates of deforestation are likely to be non-linearly related to the proportion of non-protected forest remaining; (b) land prices are likely to be influenced by the proportion of area protected in an ecoregion, and (c) the degree of uncertainty (variance) in parameter definition is likely to vary substantially among ecoregions implying that accounting for that could have a substantial impact on the solutions identified (partially acknowledged on L187). The sensitivity analyses the authors have conducted are useful for evaluating how the inferences change as the values of the fixed parameters change, but this is not the same thing as formally accounting for uncertainty in the parameters. A stronger approach would have been to specify distributions for those parameters and use robust optimisation procedures.

RESPONSE: In response to your earlier comment that begins "I agree with the perspective that conservation problems are well suited to formulation as scheduling problems," we described how we have extended our analyses to allow for dynamics of certain model components and conduct additional analyses with time-varying parameters and land costs that rise with land scarcity. Please see our response to that comment above. While we undertook a range of sensitivity analyses with varying assumptions on a range of parameters, we were hesitant to go much further on the uncertainty analysis, such as doing a robustness analysis, as we would have had to make additional assumptions about distributions, which could be influential, and therefore call into question just how robust the conclusions from this analysis would be.

(6) It assumes that climate change will not impact the distribution of forests over the specified time frame, but climate change is already altering those distributions. The authors need to explain why it is reasonable to ignore climate impacts in a long-term planning exercise.

RESPONSE: We agree that forest cover and species persistence is likely to change due to drivers other than deforestation, including climate change. As a step toward addressing this point, we have included an additional analysis leveraging projected forest cover data from the Shared Socioeconomic Pathways. To the extent that SSP or other models embed forest cover changes due to climate shifts, we can account for aggregate forest loss from multiple sources (including climate-driven shifts) in a straightforward manner. We also note that, per Figure 2, rates of forest loss (threat) play a secondary role in determining priorities, with benefit-cost ratios playing a much larger role.

Your comment does raise larger questions about other roles of climate shifts in an analysis like ours. For example, to the extent that protected areas have differential effectiveness with respect to climate-induced loss as opposed to deforestation, our analysis would need to be extended and changed. Similarly, as climate is likely to alter species ranges, the set of species in each ecoregion could be dynamic, and we would need to know how protection would alter those dynamics. Finally, impacts of climate on agriculture could alter demands for alternate land use and thus land costs (in addition to deforestation rates). In sum, climate change is likely to affect a number of aspects of conservation prioritization, some of which may have fundamental effects on prioritization.

Many of the issues raised in the prior paragraph are at the forefront of science currently and lack clear answers or consensus for how best to address them. While we acknowledge the limitations of our approach with respect to these climate-related challenges, we currently lack data to address them in a manner we would consider adequate at a global scale. Instead, we attempt to acknowledge the limitations of our approach and opportunities for improving upon it in future work in the discussion (lines 212-220):

"The type of dynamic analysis undertaken here could be expanded to include considerations of climate change and other threats beyond deforestation, parameter dynamics and uncertainty, and learning that improves conservation effectiveness through time[i]. We examined model variants that use deforestation, GDP, and budget projections derived from the Shared Socioeconomic Pathways. We find a shift away from conservation in regions projected to have rapid economic growth leading to higher relative land costs (e.g. Asia; see SI Section S3). Accounting for a full range of climate impacts would substantially increase the complexity of the problem by, e.g., introducing shifting species ranges and altering protected area effectiveness."

(7) It assumes that protected area within an ecoregion benefits all species occurring within the ecoregion. Yet ecoregions vary widely in size and, for the larger ecoregions, it is unlikely that all species would benefit equally from additional protection.

RESPONSE: We agree that, in reality, an increase in protected area is unlikely to benefit all species equally. Our extended, probabilistic approach (see Methods: Extensions to the base model) is flexible enough to account for species- and ecoregion-specific functions that relate protected area to the probability of protection for a species. The major challenge to such an approach is finding data to parameterize that function for each combination of species and ecoregion. We have added a paragraph to the discussion highlighting this and related heterogeneity (lines 235-238):

"Finally, we have not incorporated heterogeneity within ecoregions that could allow for targeted conservation in areas with greater biological benefit, lower cost, or more immediate threats. As better data becomes available, the analysis can be redone to refine our results."

Overall, all models are simplifications of reality, so it is incumbent upon the authors to make it clearer what the simplifying assumptions are and to then justify why the model strikes the right balance between tractability and relevance to the real world. I don't think that has been adequately achieved in this manuscript.

RESPONSE: Thank you for your detailed and thoughtful comments above on this topic. Where possible, we have attempted to extend our model and/or conduct supplementary analyses to address them. In remaining cases, we have attempted to make clear the limitations of our approach and the available data which would be needed to adequately address your points.

Other comments:

So little information is presented about the species data that it is not possible to evaluate whether it has been used in a suitable manner in this analysis. Providing a source citation for these data does little to justify their use. The Kier et al paper states "For 43% of the 867 ecoregions, data quality was considered good or moderate" (Kier et al 2005), which implies substantial error in richness estimates for many ecoregions. The richness estimates are central to this analysis and the authors need to summarise how those estimates were derived, justify that they are of suitable quality for this analysis, and explain why uncertainty in those estimates can be ignored (I note that none of the sensitivity analyses evaluate the richness estimates).

RESPONSE: Thank you for raising these concerns. We will attempt to separately respond to what we view as distinct points. Much of our response is in the SI, though we refer the reader to this section on lines 56-60 of the paper:

"We used plant species data because plants are an important component of overall biodiversity, plant data is available globally and is a key indicator of ecosystem productivity[20] and resilience[21]. We chose the TNC plant biodiversity data set because it has global coverage, does not focus on a particular subset of plants, and reports species richness per ecoregion."

We give a more explanation of why we thought the the TNC data (Kier et al. 2005) was the best global data set for use in our analysis in the SI Section S2:

"… for our global model, we require species richness by ecoregion for the entire set of forested ecoregions. Other sources of biodiversity data either do not have global coverage (numerous), are not attributable to ecoregions (Botanical Information and Ecological Network (BIEN)[8], Global Biodiversity Information Facility (GBIF)[9], Keil and Chase 2019[10]), and/or focus only on a subset of species (Enquist et al. 2019[11], IUCN Red List[12])."

By "not attributable to ecoregions" we mean that some sources offer data in different spatial units that cannot be cleanly mapped to ecoregions without additional assumptions of species locations within those alternate spatial units or overlap in species lists across those alternate spatial units. Deriving appropriate ecoregion-level species richness estimates from those other sources is a worthy research effort in its own right.

We agree it may be useful to the reader to understand how Kier et al. (2005) arrived at their estimates, as well as where their measures of data quality come from, and what they mean for our analysis. To this end, we have added a paragraph to the SI (section S2 under the subheading *Data Selection and Limitations*):

“More generally, we acknowledge the Kier et al. (2005) data lacks true measures of uncertainty about species richness, which limits our ability to quantify uncertainty in optimal species protection. Kier et al. (2005) arrive at species richness estimates by aggregating estimates from studies at smaller scales. The resulting ecoregion-level estimates are assigned data quality ratings, but those are based on the degree of overlap between an ecoregion of interest and the geographic area used in the source study from which Kier et al. (2005) draw. It is unclear how to translate their quality ratings to quantitative measures of uncertainty over species richness. For example, a source study could correspond to an ecoregion exactly (what Kier et al (2005) would rate as higher quality) but be based on sparse sampling and subject to high uncertainty. Conversely, a source study could cover a region having little overlap with an ecoregion but provide precise and accurate measures of richness in the target ecoregion. As such, we do not seek to use their measures of data quality directly.”

You are correct that different species richness measures would affect the optimal solution. Holding all else constant, the direction of these effects is predictable: regions for which species richness increases using an alternate data source would experience an increase in the conservation ROI at all levels of spending, which would tend to increase spending in those locations. We note that we have put in the modifier “tend to” before increase because an ecoregion may have sufficiently low conservation ROI that a small increase in the conservation ROI would still leave its ROI well below those of other regions. We believe this directional increase to be sufficiently intuitive that the real value in exploring changes to richness would come from the *magnitude* of change under *realistic* set of alternate species richness values. Unfortunately, the way in which Kier et al (2005) rate data quality does not lend itself to a quantifiable measure of uncertainty. Similarly, we lack guidance on what numbers would be appropriate for a sensitivity analysis. If you are aware of a global, unbiased measure of forest plant richness at the ecoregion scale, we would be more than happy to re-run the analysis with those richness numbers. As explained above and in the SI, we are unaware of such a source.

Finally, we fully agree that uncertainty in parameters, including species richness, is very likely to affect optimal protection strategies. As we detail in responses to other comments, properly accounting for uncertainty in model parameters and potential learning dramatically increases the complexity of the problem. The optimal solution would need to account for the potential for both passive and active learning. In the discussion, we have acknowledged treatment of uncertainty and learning as an important area for future research (lines 212-215):

“The type of dynamic analysis undertaken here could be expanded to include considerations of climate change and other threats beyond deforestation, parameter uncertainty, and learning that improves conservation effectiveness through time.”

L176. I think this inference requires qualification. There are a number of NGOs that use benefit:cost analysis (or sometimes cost-effectiveness analysis) to select among potential projects. The conclusion here that ROI is a pragmatic approximation fails to adequately explain that it still requires evaluation of all ecoregions globally in order to achieve that near-optimality. Decision makers could not use ROI on more local or regional levels (as they do now) and still expect to achieve globally optimal results. It is somewhat misleading to describe ROI as having "low data requirements" in this context.

RESPONSE: We agree our phrasing was imprecise. We have revised the text to read (lines 191-193):

“Setting conservation priorities based on the current conservation ROI may represent a pragmatic solution that achieves near-optimal results without requiring deforestation projections.”

Our intent was only to convey that data requirements are lower compared to the optimal solution, not that they are low in an absolute sense.

Planning for protection on the basis of plant species richness may not provide solutions that are good for animals. This needs to be acknowledged.

RESPONSE: We have added a caveat to the discussion (line 182):

“Protecting animal species rather than plants may, of course, lead to different solutions.”

Why is an ecoregion the best spatial unit for this analysis? (The convenience of species data being available at the ecoregion scale is not a sufficient justification). Most decisions about protected areas are made at national scales.

RESPONSE: In short, we focus on ecoregions because they are biologically relevant spatial units for which the species-area relationships (SARs) underlying our model are most likely to be applicable. Since our focus in on optimal conservation, we think biologically relevant spatial units should guide the analysis. We addressed the issue of national-level analysis in response to your first comment (see above). Even within national conservation decisions, ecoregions are often a defining unit.

Why was the problem formulated as a budget-constrained problem rather than an ecological-target constrained problem that reported the budget implications? I would argue the latter is more relevant to international policy.

RESPONSE: We agree that some conservation objectives are phrased in terms of ecological targets, especially in public communications and fundraising campaigns. However, both conservation organizations and government agencies also must grapple with the realities of budget constraints. Our interest is in adding to our collective understanding of how best to allocate those limited conservation resources, which we perceive to be the more pragmatic of the two formulations. That said, we would argue that the results we provide under varying budgets (Figure S3.1) still offer policymakers a menu of options for how many additional species could be protected at various spending levels should they wish to view it as you have suggested.

An explanation of what qualifies as a "forested ecoregion" should be added to the SI.

RESPONSE: We have added language to the SI (section S2 under the Dataset Construction) which includes how we subset the TNC ecoregion data to our 458 forest ecoregions.

"Forest ecoregions are defined as those residing in the six forest biomes as defined by Olson and Dinerstein (2002) (Tropical and Subtropical Moist Broadleaf Forest, Tropical and Subtropical Dry Broadleaf Forest, Tropical and Subtropical Coniferous Forests, Temperate Broadleaf and Mixed Forests, Temperate Coniferous Forest, Boreal Forests/Taiga, and Mediterranean Forests, Woodlands and Scrub)."

L43-47 These statements do not do justice to contemporary conservation planning. I would argue that for some time best practice in conservation planning has been to account for both opportunity and establishment/operating costs, and for planning solutions to reflect some form of ranking or prioritisation that is not formal scheduling but does at least attempt to provide information that could be used to guide the ordering of priorities in time. Moreover, stating that 'budgets stretch further in less expensive areas' is a poor way of representing the importance of ideas around maximising cost-effectiveness and return on investment. Making these sorts of weakly justified generalisations contribute nothing to establishing the merits of this work.

RESPONSE: Thank you for these suggestions. We have modified the text (lines 39-41) as follows: "Conservation strategies focus on the question of where to conserve, paying attention to areas that are biologically rich, and have high endemism or complementarity with existing protected areas[11]. Best practice conservation strategies also factor in costs[12,13] and threat[14,15]."

L73-77. Too little information is provided to understand how the species area curves were parameterised. Is A_i a measure of area or a proportion of remaining habitat? The authors must at least provide a summary of how alpha is parameterised - it is unreasonable to expect readers to dig through the supplement of another paper (ref 21) to figure this out. That reference stated: "The constant α was calculated via: (the surveyed number of species in an area of interest/(area of interest ^ z))."
and was implemented in the context of counts of parcels, not areas, so it is reasonable to expect the authors to provide some basic explanation of how the parameters were derived. The methods are not repeatable without it. Further to this, given that the parameter z appears to be part of the calculation for alpha, the authors should state more explicitly that alpha was recalculated for each of the sensitivity analyses in which z was varied (see the SM).

RESPONSE: We have revised the Methods section to make it more clear and complete. We have both changed our notation and clarified the definitions (lines 317-321) in the Methods section:

“The number of species S_{it} protected in region *i* as of time *t* is assumed to follow a species area curve $S_{it} = \alpha_i P_{it}^z$, where P_{it} is the area of forest protected and $\alpha_i > 0$ and $0 < z < 1$ govern the shape of the curve. We assume $z = 0.2$ and derive α_i as $\alpha_i = S_{i,max}/A_i^z$ in each ecoregion, where $S_{i,max}$ and A_i are the number of species and initial forest area in an ecoregion[23].”

L132-140. I thought this was a weak explanation of sensitivity to budget. A diminishing rate of return as a function of increasing budget is so commonplace that it barely warrants mentioning. A more interesting question in the context of a scheduling problem is how much more quickly species are protected as the budget increases.

RESPONSE: We agree that diminishing returns are both commonplace and a predictable consequence of the structure of the model given the concave species-area relationship (SAR). However, we also believe that examining the rate of species protection as a function of budget is subject to the same critique. Whether we examine protection over the 50-year horizon or protection in the first X periods (X=1, 2, etc), species protection will still follow the SARs across ecoregions and be subject to diminishing returns. Instead, we have opted to delete mention of declining marginal returns and simply state the quantitative species gains at different budgets. Further, because of space limitations, we have moved the discussion of results under alternative budgets to the SI. In Section S3 we include a subsection on the results under alternative budgets (and also see Figure S3.1):

“In our base model, we assumed a conservation budget of $1 billion. In reality, the global budget for forest conservation may be higher or lower than this. Information on how much is spent to conserve forest habitat on a global scale does not exist. Recent estimates of total global biodiversity finance are $78 to 91 billion annually[16], however, this figure includes spending on all ecosystems including non-forest terrestrial and marine ecosystems. Most of this amount goes to salaries and other program expenses that are separate from expenditures for land protection.

We examine a range of annual budgets from $1 million/year up to $25 billion/year. As expected, the number of species protected by the end of the planning horizon under the dynamic solution grows with the annual budget (Fig S3.1). For example, increasing the budget from $10M to $100M per year more than doubles the number of species protected (43,579 vs 20,887), as does moving from $1B to $10B per year (225,168 vs. 99,669) (Fig. S3.1).The relationship is concave, which accords with intuition given the concave species area curve relationship.”

Minor comments:

Title: This title does not adequately reflect that this analysis addresses both where and when to establish additional protected areas.

RESPONSE: Thank you for this observation. We have updated the title accordingly to “When and where to protect forests”

L6 "we identify global priorities for conserving plant species". I think this is a poor summary if what this analysis does, as it does not make it clear that this is strictly a PA planning exercise.

RESPONSE: Thank you for pointing this out. We have revised the abstract as follows (lines 3-5):

“We use dynamic optimization to identify global priorities for expanding protected areas to conserve plant species in 458 forested ecoregions over the next 50 years.”

L10-16. Isn't the dichotomy presented here artificial? In both cases the algorithm is effectively maximising ROI subject to constraints. Why devote this much space in the abstract to explaining the mechanics of the algorithm?

RESPONSE: We agree that describing the intuition before and after the first 10 years is artificial. The revised abstract instead simply describes the patterns of spending we observe in the optimal solution, in particular the concentration of spending. A key difference, which the new abstract reflects, is that initial spending focuses on “priority ecoregions” where all land is either protected or deforested, whereas in later periods, spending typically continues right up until the end of the planning horizon.

L68-69, L96 grammar / awkward phrasing

RESPONSE: We revised lines 58-60 (former lines 68-69) to read

“We chose the TNC plant biodiversity data set because it has global coverage, does not focus on a particular subset of plants, and reports species richness per ecoregion.”

Lines 77-80 (former line 96 and onward) has been revised to read:

“However, in some “priority ecoregions” with little remaining unprotected forest, relatively high species richness and low cost, the conservation ROI remains higher than in other ecoregions even after all remaining forest has been protected.”

L71. If the authors specify which year of the Kummu et al data they use I cannot find it. Mention it here?

RESPONSE: We have clarified that we use 2015 data. That is the closest year in that dataset to the most recent forest cover data available in our primary deforestation dataset (2018).

L107 suggest: spending -> allocation or prioritisation (this analysis does not describe spending)

RESPONSE: We have replaced “spending” in that sentence with “forest protection”

L108 The authors refer to "conservation gains" but I may be worth clarifying that these are actually averted losses.

RESPONSE: In the long run, yes, those gains are actually averted losses (which we agree is the point of conservation). However, in the first 10 years, they are more accurately described as protection gains. As what is now Figure 2A shows, threat plays a secondary role in determining where conservation spending goes. As a result, protected areas added in the first 10 years do not necessarily avoid species losses in the first 10 years.

L113 This is the first mention of "developed" but it is not clear to the reader what qualifies as development and how it fits into this analysis. Development may not be the best choice of word as the authors intend it to mean only deforestation, not the development of non-forested lands. I recommend replacing the word development with a more precise term throughout.

RESPONSE: Thank you for highlighting this imprecise use of language. Throughout the paper, we have removed use of the words “development” and “developed.” Instead, we exclusively use “deforestation” or “deforested” to more precisely refer to the removal of forest, which is what is relevant to our model.

L147 The qualifier 'initial' implies that the conservation benefits are not similar over the longer term?

RESPONSE: We eliminated this sentence in our rewriting to streamline the paper to make room for additional analysis and results.

L206. Species ranges do not appear to be relevant to this analysis as ranges to not appear to have been used in the Kier et al analysis to quantify richness.

RESPONSE: We have changed the word “ranges” to “richness”: “Data on species richness is also imperfect.”

L356: "and protected" is inaccurate. The authors appear to mean "and protected forest", because protected non-forested areas are not relevant to their formulation.

RESPONSE: We have revised notation for the model and now P_{it} is the area of forest protected (line 318).

L378. I don't understand where the "-1" in the exponents in this equality arise from and some other readers may struggle to understand that too. Perhaps worth explaining this more explicitly in the SI?

RESPONSE: This is an application of the power rule in differentiation to find the necessary conditions for a constrained optimization problem. The marginal benefit of an additional unit of protected area is the partial derivative of the species area relationship with respect to protected area. We have modified the text to highlight this connection:

"In this special case, the first order necessary conditions defining the optimal choices P_{iT}^{*} for all *i* for the budget constrained maximization show that the conservation return-on-investment (ROI), which measures the marginal benefit to marginal cost ratio, should be equated across any regions *i* and *j*: $\alpha_i z_i (P_{jT}^{*})^{z_i - 1} / c_i = \alpha_j z_j (P_{jT}^{*})^{z_j - 1} / c_j$."

Fig 1: There is low contrast between blue and green symbols.

RESPONSE: Thank you for raising this issue. The revised Figure 1 uses green and white symbols to increase contrast.

Comments on Supplemental Information:

p2 bullet point, typo "iin", and "we use the species-area relationship to derive iin each ecoregion" does not make sense. To derive alpha_i in each ecoregion? If so, please provide more explicit detail on how the alphas were derived and what their values are. Why not report those alpha values in Table S5, which could be reduced in font size to fit on the page?

RESPONSE: Thank you for identifying the typo, which has been fixed. Regarding α, we show how we derive α in the Methods section from a simple rearrangement of the species-area relationship (lines 319-321):

"We assume $z = 0.2$ and derive α_i as $\alpha_i = S_{i,max}/A_i^{z}$ in each ecoregion, where $S_{i,max}$ and A_i are the number of species and initial forest area in an ecoregion[23]."

Per your suggestion, we also now add α values for our base scenario to what is now Table S3 in the SI (formerly Table S5).

Providing 11 maps is good, but readers would benefit from a matrix of statistics that helps to evaluate the similarities or differences among solutions. It is almost impossible to gauge this objectively by flipping between pages to compare maps.

RESPONSE: Thanks for this suggestion. We have added a new Table S2 to the SI with correlations in optimal species protection vectors (one entry per ecoregion) using the dynamically optimal solution under different assumptions as compared to the base case. Correlations are presented both for protection in the first 10 years and the full 50-year planning horizon. In the first 10 years, all correlations are quite high (above 0.94) except one. That lowest correlation (0.74) arises under the alternate cost assumption and is to be expected because the correlation between our alternate costs and base costs is only 0.79. Given the central role that costs play in conservation ROI, it is unsurprising that such a dramatic change in costs would alter the optimal solution. Expanding to the first 50 years, correlations of protected species vectors are as high or higher than correlations for the first 10 years, with the exception of longer planning horizons (100, 200, and 725 years). Those lower (but still high) correlations for longer planning horizons arise because the optimal solution only uniquely defines protection at the end of the planning horizon. For a longer horizon, optimal protection after 50 years is not uniquely defined; an alternate protection trajectory that reaches the same final protection levels but has higher correlation after 50 years may exist.

REFEREE #3
(Remarks to the Author):

We thank the referee for the excellent set of comments, which pushed us to think harder and go further in the analysis. We think the changes motivated by your comments and those of the other reviewers, have significantly improved the paper. We detail how we responded to your comments below (our responses in blue).

Referee report on “Where to protect first in global forest conservation”
by Ian Luby, Steve J. Miller and Stephen Polasky

Overview

The paper focuses on the issue of how we make decisions regarding conservation. To my mind this is a considerably more important research frontier than that addressed in the bulk of conservation research papers.

That said my first reaction was that this might be a somewhat synthetic comparison of two or more academically concocted approaches. I was therefore relieved to see that the authors contrast these with some of the most commonly adopted approaches to decision making. Adding to this their finding that the use of straightforward and relatively available heuristics provide reasonable approximations to optimised solutions and I feel that this paper is both a strong contribution to the literature and a useful guide for practical policy implementation.

Essentially the paper is concerned not with what to conserve but rather when to conserve it; showing that temporal ordering is of crucial importance in a world of rapidly advancing species loss. This provides something that I always enjoy in a paper – the realisation of what in hindsight looks immediately obvious and clearly important; it’s a nice result and an important contribution to the literature.

RESPONSE: We thank the referee for the kind words and the encouragement. We also thank the referee for the excellent set of comments that follow.

Specific comments

Line # Comment
39-41 Please support the following statement with a reference:

“Conserved lands are not well chosen to maximize biodiversity conservation but consist in part of land that is low priority for human uses.”

RESPONSE: We have revised the language slightly (“Moreover, conserved lands are not always chosen to maximize biodiversity conservation but consist in part of land that is low priority for competing human uses.[10]”) and added the following reference:

Joppa L.N., & Pfaff, A. High and Far: Biases in the Location of Protected Areas. *PLoS ONE* **4**, e8273 (2009).

70-71 “Land cost was determined using globally gridded GDP from Kummu et al.”

I had to go and read the Kummu et al. paper to try and make sense of this approach as my initial reaction was that it seemed a very strong assumption that the GDP of a country, divided by its total land area, should give an unbiased estimation of the per hectare cost of land for conservation in that country. I tried to think of examples where that was clearly true but instead find lots of examples of where this might be quite erroneous, especially for small prosperous countries: Singapore is relatively tiny but has a huge per capita GDP implying a high per hectare GDP; Macao, Luxembourg and Brunei are similar examples. The opposite might apply in very large, low GDP countries.

I admit that, looking at the Kummu et al. paper the maps do give a sensible looking spatial distribution – but the assertion that this validly relates to per hectare conservation costs, especially in rural areas still seems one that deserves more justification. This justification is not forthcoming in the Supplementary Information.

Presumably two countries with similar population and GDP but one of which has an IT dominated economy while the other excels in high value agriculture would be assumed to have similar land costs. That seems unlikely to be the case in practice.

My major concern here is that any error arising from this assumption might not be unbiased. Won’t this lead to net benefit measures being biased towards very large countries with low population/GDP per capita?

The assertion that land cost is itself a proxy for conservation cost is also something that needs defence.

I feel this approach needs further discussion and justification, at least in the SI.

RESPONSE: Getting globally consistent estimates of land cost is a challenge. We choose the Kummu et al. data set because it provides a globally consistent, complete, and recent data set of the value of economic activity (GDP) at a grid cell level. The Kummu et al. data takes the most spatially disaggregated information about GDP per capita, which can be state or provincial level data or national data, and uses gridded population data to get estimates of GDP by grid cell, with higher population getting higher GDP.

We provide a brief description of the Kummu et al. data set in the SI (Section S2):

“Land cost was measured using globally gridded GDP data from Kummu et al. (2018)[5]. Specifically, we use the average per-capita GDP produced per unit of forested area in an ecoregion as our proxy for land cost. Given the gridded nature of available GDP data and forest cover data, we must make an assumption about how to treat GDP in partially forested pixels. In the main text, we assume that GDP within a pixel is distributed uniformly, so that in partially forested pixels we multiply GDP by the fraction of the pixel covered by forest. We consider an alternative assumption for treating partially forested pixels in Section S3 below."

We have also included language in the SI (Section S3) about why we chose the Kummu et al. data set over other datasets:

"Global land cost data provides another challenge. We selected the Kummu et al. (2018) data set because it is globally consistent, complete, and recent. The data consists of spatially disaggregated values of economic activity that used widely available national and sub-national GDP and population data for the entire world, allowing us to use it for our 458 ecoregions. Other global datasets are limited in the type of land they cover (Naidoo and Iwamura 2007)[13] or are not as complete or recent as Kummu et al (e.g. G-Econ dataset[14], UNEP/GRID-Geneva[1]). As we acknowledge in our paper, using GDP-based metrics as a proxy for land cost has its drawbacks, but it is the best proxy available for our model."

We acknowledge there are remaining issues with using gridded GDP data for our purposes. First, we assume that the measure of economic activity (GDP) is correlated with land costs. Areas with high value of economic activity, either because they are more densely populated or tend to be richer, will generally have higher land values. We think that using gridded GDP as a proxy for land cost is reasonable, though it is clearly not a perfect proxy.

We note several issues with use of this data and with land costs more generally. We describe some of these issues in the Discussion (lines 232-235):

"Issues with gridded GDP include being an incorrect signal of the value of forest land if the value of economic activity is not primarily related to forests or to land use, significant subsistence or cultural use values not reflected in GDP, and not incorporating maintenance and enforcement costs."

In reviewing the Kummu et al. data set for ecoregions that we know, we found a good match between the data and what we know about relative land costs in these regions. Granted this is a somewhat intuitive and informal assessment. However, it does lead us to believe that the Kummu et al. data does not generate a systematic bias that would affect results in important ways.

One thing that we found in doing this work is the need for better proxies for conservation cost that are globally comprehensive and consistent We are working on such a data set in another project but this work is not yet complete and so cannot be used for this current analysis. We agree that Kummu et al. provides only a proxy for land cost—let alone conservation cost—and have tried in the text and the SI to be clear about its limitations.

74 The species–area relationship (SAR) adopts a classic power function form as per Preston (1962). This transformed linear approach is widely adopted in the literature, but has been the subject of serious critique (see Leitner and Rosenzweig, 1997; Lomolino, 2000) and is fundamentally challenged by non-linear, tipping point models which increasingly have empirical support. My own view is that the power function form is acceptable when we are considering changes which do not threaten species loss. However as that exception is precisely the focus of much contemporary research including this paper then I feel that alternative perspectives should at least be acknowledged. I am not asking that the analysis should be reconsidered; the prevalence of the power function SAR means that it is a reasonable approach and this is not the novel contribution of the paper (and that contribution is very important). But some acknowledgement of this issue would be welcome.

RESPONSE: Thank you for raising this important issue. We agree with your concern and we appreciate you supplying references. In the Discussion, we've added references to Drake and Griffen (2010) and Lomolino (2000) along with an acknowledgement of this issue. This acknowledgement is included with other concerns pertaining to our knowledge of species and their relationship with area:

"Other issues that could alter species-area relationships include potential tipping points and changes in patterns for small populations[31,32], quality of forest habitat, degree of fragmentation and landscape-level integrity[28,33], and survival of forest species in non-forest habitat, all of which could alter species-area relationships."

References:

Drake, J. and Griffen, B. (2010) Early warning signals of extinction in deteriorating environments. Nature, 467, 456–459 https://doi.org/10.1038/nature09389

Lomolino, M.V. (2000), Ecology's most general, yet protean pattern: the species-area relationship. Journal of Biogeography, 27: 17-26. https://doi.org/10.1046/j.1365-2699.2000.00377.x

Leitner, W.A. and Rosenzweig, M.L. (1997) Nested Species-Area Curves and Stochastic Sampling: A New Theory, Oikos, Vol. 79, No. 3, pp. 503-512 https://www.jstor.org/stable/3546894

Preston, F. W. (1962) The Canonical Distribution of Commonness and Rarity: Part I, Ecology, Vol. 43, No. 2 pp. 185-215

76 The preceding lines discuss the relationship between species conserved and the area of remaining habitat; the species–area relationship (SAR). A line in the paper that jumped out for me was regarding the curvature of this relationship given by parameter z where the authors write "We assume $z = 0.2$". Although reference is made to Wilson et al (2006), this still seemed a pretty uncompromising statement, so I was pleased to see in the SI that the authors examine

a slightly wider set of values for z ranging between 0.1 and 0.3. Nevertheless, even this felt like it needed further justification and so I took a look through Wilson et al (2006) and found it to be similarly opaque in the main paper. The Supplementary Methods to the latter paper simply states that "We employed a z of 0.2, however our results are insensitive to the value of z.". So again the value of $z = 0.2$ is asserted rather than empirically based.

Subsequent research has shown that the Wilson et al (2006) assumption of a single SAR for all species can be rather simplistic. Matthews et al (2014) show that generalist and specialist species exhibit significantly different SARs. Given that the proportion of generalist to specialist species is unlikely to be constant globally then assuming a single SAR is at best a simplification which needs to be highlighted. I don't think this is a 'killer problem' but it does mean that readers should be somewhat circumspect about treating the detailed results in the SI as definitive. This however does not take away from the very real contribution which the paper makes in terms of highlighting the temporal ordering problem – that is an important issue. Rather I just feel that the authors need to acknowledge that future research should address these real world complexities.

RESPONSE: We appreciate this comment and as a result have worked to explore this issue more thoroughly. In the SI, we include additional analyses taking different values for z. In the initial submission we had uniformly varied the z value to be z = 0.1 and z = 0.3 in addition to z = 0.2. As the referee rightly points out though, there are reasons for thinking that z values might differ by ecoregion. We did two additional analyses with different z values by biome, from Kier et al 2005, and from Storch et al. (2012), which we describe in the SI (Section S3):

"The change in the conservation ROI with additional protection depends on the curvature of the species-area relationship in a given ecoregion. In the main text, we assume $z = 0.2$ across all ecoregions, as in Wilson et al. (2006). Here, we examine four alternatives: $z = 0.1$ (Fig S3.2), $z = 0.3$ (Fig S3.3), z drawn from Kier et al. (2005) (Fig S3.4), and z drawn from Storch, Keil, & Jetz (2012)[17] (Fig S3.5). In that final case, we use values in the SAR slope ("lower half") column from Table 1 of Storch, Keil, & Jetz (2012) since forested areas in all of our ecoregions fall within the range (0 to $10^{6.1}$ km^2) defining applicability of that column. Ecoregions are mapped to continents based on maximal areal overlap; those not falling in the continents in Storch, Keil, & Jetz (2012) are assigned the average of the SAR slope ("lower half") values in that study. While some differences in optimal species protection arise, broad patterns of priority regions of the globe remain stable."

Reference:

Matthews, T.J., Cottee-Jones, H.E. and Whittaker, R.J. (2014), Habitat fragmentation and the species–area relationship: a focus on total species richness obscures the impact of habitat loss on habitat specialists. Diversity Distrib., 20: 1136-1146. https://doi.org/10.1111/ddi.12227

83-91 [my comment re lines 132-138 supersedes some of this – making clear at 83-91 that you examine sensitivity to budget changes would remove much of the issue here]

The choice of budget might appear at first to be a rather trivial issue; surely whether it is $1Billion or $100Billion is irrelevant to the fundamental findings. However, I doubt this is the case; the bigger the budget, the more can be done now and so the less that the negative impact of sub-optimal timing matters. If we had sufficient budget to address all conservation needs this year then the timing issue effectively disappears. Conversely, the lower the annual budget the greater the potential for temporal problems; and the chosen budget is on the low side and arguably might have been selected to highlight the issue.

Now the simple fact is that the hypothetical case above will remain so – this is simply not going to happen and the temporal effect is real. But I feel that it is worth acknowledging the relationship between budget and temporal sub-optimality.

RESPONSE: We agree that more budget is better for conservation and if the budget was high enough, we could protect everything today and there would be no need to assess dynamic tradeoffs.

In line with your suggestion, we now mention that we examine sensitivity to budgets earlier (lines 71-72). As the results using different budgets are as expected given the shape of the species-area curve, and given space constraints, we decided to move the budget sensitivity analysis itself to the SI. This analysis is described in the SI Section 3 (which we reference in lines 71-72 and lines 127-128 in the text) and is shown graphically in Figure S3.1. We have rewritten a paragraph in the Results section to discuss budget-related issues (lines 127-133):

“We analyzed the benefits of increasing the budget (which, following the shape of the species-area relationship, has diminishing returns; see SI Section 3 and Fig. S3.1) and the benefit of spending the 50-year budget up front. The gains from immediate access to the entire budget are modest, an additional 124 species (0.12%). Budget size and timing matter to the extent that it allows conservation to proceed before deforestation, but allocating the annual budget to priority ecoregions captures much of this potential, leaving little remaining gain from shifting spending more earlier in time.”

92-100 I felt that the description of the rules used for dynamic optimisation was a little too brief to aid clear comprehension by the reader. What is the unit used here? Is it the “species-dollar” per hectare value suggested at line 78? i.e. are the authors ranking every single separate hectare of potential conservation land globally and ticking them off sequentially so that the first hectare might be in Sumatra while the second is in Madagascar and so on. Or are they taking this on an ecoregion by ecoregion basis? Or is some other rule being used? It wasn’t clear to me and this matters because each approach will lead to different outcomes (and different levels of institutional cost and feasibility).

This just needs clarification.

RESPONSE: Thank you for this comment. We realize that our initial description of the model and the optimization method was insufficiently clear. We have substantially revised the description of the model and the optimization. We have revised the description in the main text (lines 63-72) to give a short and more intuitive explanation of the model, and greatly expanded and revised the Methods section (lines 315-487). The Methods section now begins by defining variables used in the basic model and formally defining the optimization problem (lines 312-340). We then describe the intuition for the optimal solution (lines 341-361). We describe our approach to numerically solving for optimal solutions (line 363-394) and extensions to the basic model (lines 396-475). We conclude the Methods section with a description of solutions using heuristic approaches (lines 477-487).

In each period, the conservation planner has the opportunity to increase protected area in any ecoregion, subject to both budget and feasibility constraints. In your framing, this can be viewed as ranking (infinitesimally small) units of area across all ecoregions and purchasing the unit with the highest conservation ROI, regardless of which ecoregion it falls in. The units must be re-ranked after each allocation owing to the interdependencies in marginal value due the nonlinearity of the species-area curve. The dynamic nature of the problem also means that ranking must be based on the *final* conservation ROI when all opportunities for further protection have been exhausted (by the end of the planning horizon or because no forest is left to protect).

In the text, we have attempted to resolve your question with two specific phrases to emphasize that we do not impose spending solely in specific ecoregions by ranking ecoregions:

- Lines 66-67: “Each year, currently unprotected forest in *each ecoregion* can be protected via purchase…” (emphasis added)
- Lines 82-83: “This initial allocation of all budget to a subset of ecoregions is an outcome of optimization (see SI Fig. S1 for intuition).” We discuss this point further in the SI (Section 1).

103-104 This states:

“In the first 10 years, with a $1 billion annual budget, the optimal solution involves selecting habitat conservation in 178 out of 458 forested ecoregions”

This is an interesting statement.

First, it goes some considerable way towards answering my previous point (re lines 92-100) but I am going to leave that comment in because I feel it shouldn’t be down to the reader to works these issues out – and also it begs the question; why do it this way? Why look at entire ecoregions sequentially, solving the problem there while doing nothing elsewhere? I am not saying that is the incorrect approach – I just feel that the issue and potential trade-offs should be clarified for readers.

RESPONSE: Per our response to your comment on lines 92-100 above, we rank potential allocations of the marginal dollar. The fact that the optimal solution allocates all money to a subset of ecoregions is not a choice or a restriction imposed upon the optimal solution. Instead, it is a consequence of the biological and economic features of the data. A small number of ecoregions have higher conservation ROI than other ecoregions, even if we protect all remaining forest in those ecoregions. As a result, the repeated allocation of the marginal conservation dollar will go back to these regions in the first several years not by design, but because that is where we get the most (species) for our money. Because this focus on a small number of ecoregions is the result of optimization, we feel that renders a discussion of tradeoffs unnecessary (but we'd be happy to revisit if you still feel this is necessary even after our approach has been clarified).

Second, the selection of a 10-year initial period is somewhat arbitrary. Looking at the results reported in Figure S3.1 of the Supplementary Information it seems very likely to me that the initial years of intervention (corresponding to the initial budget allocation) have much greater marginal impact on conservation than subsequent years (indeed this is the major thesis and contribution of the paper). Might it not be nice to make the most of this by reporting for more than just the two periods (up to ten years and beyond ten years) shown here? If not in the main paper then this might be a useful illustration for the Supplementary Information which would really bring home to reader (including, hopefully, policy makers) then need for urgent action now. Such an approach might also overcome the visual problem raised by the fact that, as lines 110-112 state, "most ecoregions which eventually are targeted for conservation over the entire 50-year planning horizon are selected within
the first 10 years (178 of 205)." Visually this makes those areas targeted after the first ten-year period difficult to discern at first.

RESPONSE: Thank you for this comment. It has made us rethink how we discuss the time component of optimal protected area expansion.

First, we believe the need for urgent action *is* emphasized by a focus on the first ten years. In fact, the first paragraph of our results emphasizes the large fraction (44.1%) of species that are protected within the first years (as compared to those protected after 50 years). That said, we would argue that is less a feature of our solution and more a consequence of the species area relationships. Those relationships exhibit declining marginal returns, so that the first dollars spent in any ecoregion will provide the biggest boost to species conservation. Even a sub-optimal approach that ignored the dynamics of deforestation would thus likely deliver the highest conservation benefits in the first several years. So, while we do emphasize the early gains in species protection, we do not think they are sufficiently unique to our solution to place even more weight on them.

We also note that the time path of aggregate species protection under our optimal approach is represented in what is now Figure 2, panel B. The purple solid line shows the accumulation of species protection through time. While in general this also emphasizes the importance of protection in initial years, it also adds nuance. As stated in the discussion, the optimal solution does actually shift to slightly lower conservation ROI areas than a pure benefit-cost heuristic does. That temporary dip protects areas that would later be lost to deforestation, with the optimal solution later returning to higher conservation ROI areas and surpassing the benefit-cost heuristic.

Finally, we have taken to heart your suggestion about how the focus on the first 10 years in Figure 1 made subsequent protection difficult to discern. In response, we have changed Figure 1 to depict protection over the entire 50-year horizon (circle size), instead using colors to group ecoregions based on what fraction of that final protection occurs in the first 10 years. While we agree that 10 years is still an arbitrary cutoff, it does facilitate and simplify explanation. In conjunction with the accumulation curve in Figure 2, we hope this provides the reader with a more complete picture of how conservation budget allocation plays out across both space and time.

106-108 The statement:

“There is also a conspicuous lack of spending in many of the richest regions of the world including the USA, Europe, Japan and Eastern China.”

could easily be misconstrued. It’s not that these countries are not or should not be spending on conservation; rather it is that the geographical focus of their spending would be best directed at other areas of the world. Indeed, if they really are motivated by the needs of wild species rather than domestic use values or national political gains then their funding should be directed overseas in this initial period as it will deliver much greater improvements in global biodiversity conservation.

RESPONSE: We thank you for raising this important distinction. We have changed the word “spending” to “forest protection” to emphasize that what we measure is protection, not the source of funds (which, in our model, are not associated with a particular national source). That sentence now reads:

“In contrast, little forest is protected in many higher income regions with high land costs including the USA, Europe, Japan and Eastern China.”

113-114 The statement “In 124 of these 178 ecoregions, all land is either conserved or developed by the end of the first 10 years” puzzled me. The term development has not been mentioned before. Can you clarify your meaning here; is there something I have missed in the description of the analysis?

RESPONSE: Thank you for highlighting this imprecise use of language. Throughout the paper, we have removed use of the words “development” and “developed.” Instead, we exclusively use “deforestation” or “deforested” to more precisely refer to the removal of forest, which is what is relevant to our model.

122-123 I do not understand the term “largely” in the following:

“Spatial patterns of early species protection are driven largely by initial conservation return-on-investment (ROI),” What else is driving results apart from ROI? How is Threat (in Fig 2) defined? Is it brought in to the analysis and if so then how?

RESPONSE: In principle, both threat and conservation ROI determine optimal conservation patterns. However, our analysis of the ecoregion data showed that conservation ROI was much more important than threat. We have rewritten our explanation of the basic results to emphasize this point (lines 117-126):

“Spatial patterns of early species protection are most strongly associated with high initial conservation ROI, while there is little correlation between the rate of deforestation (threat) and species protection (Fig. 2A). The optimal solution seeks to equalize final conservation ROI, which favors prioritizing ecoregions with high initial conservation ROI. Threat could lead to protection in ecoregions with lower initial conservation ROI if the ending ROI is at least as high as in other areas (see SI Figure S1 for an example). This is more likely to occur in ecoregions with little remaining unprotected forest or with species-area curves with little curvature. However, because this combination rarely occurs in our data, we find only a modest role of threat in our application. SI Section S1 further discusses the limited role of threat.”

Similarly, the Figure 2 caption now reads “Species protection is high in the first 10 years where the initial conservation ROI is high (ecoregions in the upper left). Level of threat is not highly correlated with initial spending.” We have also clarified the definition of threat in the caption of Figure 2: “Threat is defined as the fraction of remaining unprotected forest that is deforested per year at current rates.” While threat is not explicitly incorporated in that way in the analysis, protection is constrained by remaining available forest. The higher that measure of threat, the more likely deforestation is to constrain optimal protection.

127 - 129 “fast-moving development also limits potential species conservation by reducing the remaining forest area that could be protected.”

It is not clear that you capture spatially different speeds of development in your analysis – so are you really referring to the speed of conservation and implicitly assuming that development proceeds at the same rate globally?

RESPONSE: We capture different deforestation rates for each ecoregion. We calculate deforestation rates based on the rate of loss of forest area between 2000 and 2018, which we describe in the main text in lines 52-54:

“For each forest ecoregion, we used the area of forest remaining in 2000 and 2018 and the annual rate of deforestation implied by those data…”

In the updated model notation shown in the Methods section, this is captured by the deforestion rate d_{it} having subscript i. In our base case analysis we assume that deforestation rates are time-invariant within each ecoregion but vary across ecoregions. We also considered time varying deforestation rates. In the SI, we examine alternate assumptions in which deforestation rates also change through time following Shared Socioeconomic Pathways (SSPs) (SI Section 3).

132-138 Great to see the budget sensitivity analysis here.

It seems to me that there could be an interesting trade-off analysis here, working out the cost-savings afforded by adopting a temporal optimisation approach, as follows:

• Suppose you consider a world where we ignore temporal optimisation (i.e. the present situation) and we compared one low conservation budget $X1 (delivering conservation x1) with a higher budget $Y1 (delivering conservation y1). This leads to an improvement in conservation of y1 – x1 = z1.

• Now consider a world where the conservation budget is fixed at $X but we can now switch on temporal optimisation. This leads to a new level of conservation x2 where x2 > x1.

• Let’s suppose that x2 < y1 (it doesn’t actually matter what this relationship is as long as x1 > x2, which you have already proven to be the case)

• Now increase the budget from $X1 to $X2 and, using temporal optimisation until x2 = y1 i.e. the conservation improvement is again z1

• The use of temporal optimisation should mean that $X2 < $Y1

• The difference $Y1 - $X2 should give us the value of information provided by temporal optimisation.

By varying the initial budget amount $X2 you could get a relationship between budget and conservation effectiveness under temporal optimisation which might provide a useful policy message.

Showing policymakers that, by using existing data we can very substantially improve the value for money delivered by conservation spending has to be one of the strongest findings of the paper.

RESPONSE: Thank you for this interesting and powerful alternative framing of the value of dynamic optimization. We’ve opted to include this cost savings perspective as an additional way to understand our comparison to heuristics. We find the heuristics to be an appropriate set of comparisons that all ignore dynamics in various ways (along with other aspects of the problem). As we have already examined optimal species protection as a function of conservation budget, we can compare species protection under heuristic approaches to optimal protection at various budget levels to put a lower bound on cost savings due to optimization. In other words, we can find the smallest budget for which the optimal solution conserves at least as many species as the heuristic approach, which provides a lower bound on cost savings. Using this approach, we find that, with the exception of the benefit-cost heuristic, which performs near optimally, the dynamically optimal approach offers cost savings of at least 75% compared to other heuristics. We have included this comparison in the paragraph comparing the optimal and heuristic solutions (lines 143-146):

“Viewed from a cost savings perspective, the dynamically optimal approach offers at least 75% reduction in costs for the same levels of species protection compared to the heuristics, with the exception of the conservation ROI heuristic (see SI Section S5.1).”

SI Section S5.1 explains the calculation of these cost savings and provides details on each heuristic.

149 – 150 I really like the idea of using heuristics. The approach seems resonant of Wilson et al (2006) who might be worth referencing here as well as at the end of the section (as they are at line 159)

RESPONSE: Thanks for the suggestion. We have included the a reference to Wilson et al. (2006) here.

151 – 158 Great results which have a strong policy message.

RESPONSE: Thank you.

170 – 175 “The highest priority forest ecoregions for conservation globally lie in Melanesia, Southeast Asia, China, the Anatolian Peninsula, the Amazon, and Central America”

I suggest that you here reinforce the message noted in my comments on lines 106-108 above.

RESPONSE: We agree with your comment. We have tried to emphasize this point by placing it several places in the main text. We refer to the location of priority ecoregions in the second sentence of the Results section (lines 102-104):

“The largest gains in species conservation come in Melanesia, Southeast Asia, parts of China, the Anatolian Peninsula, northern South America, and Central America (Fig. 1).”

We also refer to these priority ecoregions in the legend of Figure 1.

We return to the location of priority ecoregions in the opening paragraph of the Discussion section (lines 171-174):

“Over the planning horizon, conservation spending occurs in 45% of forested ecoregions (205 out of 458) and is concentrated in Melanesia, Southeast Asia, parts of China, the Anatolian Peninsula, northern South America, and Central America.”

Discussion section This raises a number of caveats (some of which I raise above). However, I feel it could usefully be more positive about the importance of the central message of the paper – that the question of when we should conserve is a vital and under-researched issue which turns out to be of great importance to the challenge of addressing biodiversity loss.

RESPONSE: Thank you for this helpful comment. When starting in on our revisions and being diligent about responding to each referee comment, we seemed to be adding more caveats and more words until the Discussion section became too disjointed, too long, and too negative. We then followed your suggestion and streamlined the Discussion, putting some discussion of alternative assumptions in the SI and rewriting the Discussion to be more positive in tone and more unified in approach.

The advantages of countries agreeing to act cooperatively in ways which yield the best returns on available resources is also worth stressing.

RESPONSE: Thank you for this suggestion. We agree, which is one of the reasons we did a global analysis that would result if countries could in fact cooperate to optimize species protection globally. To address this and your following comment, we have added the following in the Discussion (lines 198-204):

"While our approach is equally applicable at smaller scales (e.g., national protected area planning), we conduct a global prioritization to emphasize potential species protection gains with global cooperation. When conservation planning is done at national scales, it is likely there will be a mismatch between conservation opportunities and budgets leading to unequal conservation ROI across regions, generating lower protection than does coordinated global planning[25]."

Somewhat contradicting this - but perhaps with a more realistic view of feasible policy action - is it worth adding that the issue of temporal optimisation is relevant even at the national level.

RESPONSE: Thank you for this suggestion. We agree. See prior response.

Finally, I also felt that the brief discussion of Figure 3 within this section should be extended to reinforce the message that simple heuristics such as net benefit analysis can help realty in prioritising action.

RESPONSE: Thank you for this comment. We agree that the result that the benefit/cost (conservation ROI) heuristic does so well is an important finding of the paper. We now mention this finding in the abstract (lines 11-12):

"A strategy using conservation return-on-investment while ignoring threat is close to optimal…"

We discuss this finding in the Results section (lines 135-139)

“A heuristic that invests in ecoregions with the highest current conservation ROI (marginal benefit to marginal cost ratio) performs remarkably well, conserving 99.45% of the species protected by our optimal approach. These results are similar to findings in Wilson et al. [Error! Bookmark not defined.] and highlight both the limited role of threat and the potential benefits of well-designed heuristics.”

In the Discussion, we provide more extensive comments on this finding and what we view as the significance for providing pragmatic guidance (lines 183-195):

“We find that choosing to conserve in ecoregions with the highest current conservation ROI achieves over 99% of the species protection of the dynamic optimal solution, but does not require data on deforestation rates, remaining forested area, or complex dynamic calculations. This heuristic performance is especially encouraging given that the dynamic optimal solution protects fewer species early on, catching up only later on (Fig. 2B). While that temporal tradeoff is optimal, locking in moderately high conservation ROI where it would not be possible at a later date (see Area C in SI Fig. S1), it carries risk. If future rates of deforestation or future budgets are uncertain, later gains may not be realized. Setting conservation priorities based on the current conservation ROI may represent a pragmatic solution that achieves near-optimal results without requiring deforestation projections. In other contexts, explicit consideration of dynamics may play a larger role, especially if more places with moderately high conservation ROI face immediate threats and have little remaining habitat.”

REFEREE #4
(Remarks to the Author):

Thank you for your positive comments. It always boosts spirits to have a referee say "This is an excellent article…" We appreciate it. We also thank you for the comment about the importance of climate change. We think the changes motivated by your comments and those of the other reviewers have significantly improved the paper. We detail how we responded to your comments below (our responses in blue).

This is an excellent article that will be a very useful contribution to the ongoing effort to find ways to protect biodiversity worldwide. The conception and statistical modeling are original and will be of general interest. The conclusions are robust and well presented so that others can analyze these data in comparison well. The references are full and adequate, and the writing is clear.
I suggest that they add a short paragraph or even a few sentences about climate change. In my own studies, as of the disappearance of moths and butterflies, it is pretty clear that in the quite near future, climate change will become an even more important driver of species and population loss than land use, which of course has been dominant over thousands of years.

RESPONSE: We thank you for the suggestion to bring in some discussion of climate change impacts. We agree about the importance of climate change considerations in conservation. We have incorporated some aspects of climate change in our analysis by using Shared Socioeconomic Pathways and associated Representative Concentration Pathways (SSP/RCP) projections for forest cover and GDP per capita, which is then also used to change land costs. To fully address climate change would require going much further in the analysis, including thinking about how climate would shift species ranges and forest habitats. Many of the issues raised in with climate change are at the forefront of science currently and lack clear answers or consensus for how best to address them. We attempt to acknowledge the importance of climate change, the limitations of our approach, and opportunities for improving upon it in future work in the Discussion (lines 212-220):

"The type of dynamic analysis undertaken here could be expanded to include considerations of climate change and other threats beyond deforestation, parameter dynamics and uncertainty, and learning that improves conservation effectiveness through time[27]. We examined model variants that use deforestation, GDP, and budget projections derived from the Shared Socioeconomic Pathways. We find a shift away from conservation in regions projected to have rapid economic growth leading to higher relative land costs (e.g. Asia; see SI Section S3). Accounting for a full range of climate impacts would substantially increase the complexity of the problem by, e.g., introducing shifting species ranges and altering protected area effectiveness."

REFEREE #5
(Remarks to the Author):

We thank the referee for the excellent set of comments, which pushed us to think harder and go further in the analysis. We think the changes motivated by your comments and those of the other reviewers, have significantly improved the paper. We detail how we responded to your comments below (our responses in blue).

General Comments

Luby et al. identify conservation priorities for forested ecoregions globally based on information for each of the 400+ forested ecoregions: plant species richness, estimated cost of conservation, degree of protection, and threat. The advance over other recent work is that they do this together with an optimization that considers the sequence of protection through time.

The biggest gains would be in the first 10 years (half of the total species protected) with more moderate gains in the 50 year time horizon as less efficient options are available.

This is an important paper that begins to address the conservation actions in sequence. I think this complementary perspective is useful in addition to the biodiversity analyses based on spatial priorities and irreplaceability, although I also have my reservations about this explained below. I do appreciate all of the sensitivity analyses and the discussion of the limitations of the study. The focus on plants and forested ecoregions is also a nice addition to the studies focused on vertebrates, and certainly warranted from a conservation perspective.

RESPONSE: Thank you for seeing the value in our approach and results.

One of my main reservations is the use of species richness/protection at an ecoregion level. For example, the metric 'species protected in first 10 years'. But we don't actually know which species are already protected, and therefore, don't know which species need to be protected in the future. If we have two ecoregions with the same species richness and the same amount of forest in protection, but one ecoregion has reserves in high quality habitat and the other has reserves in species-poor regions, wouldn't these very different scenarios be effectively the same thing for the purposes here? I realize that changing the z value on the species-area curve is investigated in the appendix, but it doesn't seem to address this problem. This limitation is discussed (Lines 219) and I realize that spatially explicit data is not high enough quality for many plant species to be used systematically, but are there species lists from protected areas ? Could GBIF or similar be used to generate lists inside and outside of protection within each ecoregion? Or perhaps GBIF could be used to generate an estimate of overall sampling effort in each bioregion to be factored into the species-area analysis?

RESPONSE: You raise at least two interesting and related points with this comment. The first is that protection may fail an additionality criterion by failing to protect new species. Second, our protection efforts are subject to uncertainty in that we do not know exactly which species will be helped by additional protected area.

In response to the first point, we appeal to the fact that we are solving a global conservation problem with a goal of highlighting possibilities: how many species we *could* protect for a given budget. Several of our assumptions are in keeping with this perspective. The global budget is appropriate only if there is international cooperation, and we assume that protection is perfectly effective. By assuming a specific protected species-protected area relationship, we are implicitly making another optimistic assumption: the highest marginal value areas of forest will be protected first. In short, we're assuming that once we allocate conservation budget to an ecoregion, local experts will allocate that budget such that the best areas will be protected first, hence tracing out the species area curve.

While we argue this view is defensible given our optimization goal, we acknowledge both here and in the discussion that heterogeneity -- in habitat quality and elsewhere -- is very likely to be present (lines 235-238):

"Finally, we have not incorporated heterogeneity within ecoregions that could allow for targeted conservation in areas with greater biological benefit, lower cost, or more immediate threats. As better data becomes available, the analysis can be redone to refine our results."

We considered the use of GBIF data for the purpose you suggest and others. After extensive investigation and discussion with colleagues, we decided that using GBIF data would introduce more problems than it would solve. In particular, we think that because sampling effort is not only different, but sufficiently biased both across and within ecoregions that we would not get useful results about conservation priorities using GBIF data. We didn't use information from protected areas because we were committed to having comprehensive coverage of all forest ecoregions.

We also acknowledge the validity of your second point: not only could there be suboptimal allocation of conservation budgets within ecoregions, but there could be substantial uncertainty about the marginal benefits of increasing protection in a particular location. One approach would be to reduce uncertainty with better spatially resolved data. Per the comments above, we do not believe a suitable dataset currently exists, at least not for plants. Another approach would be to explicitly account for uncertainty in marginal benefits of conservation in the optimization problem itself. However, this would dramatically increase the complexity of the problem, requiring considerations (and additional, likely unpalatable assumptions) about the precision of new information, and the potential for learning to reduce uncertainty.

My other main reservation was that the paper overlooked some fundamental aspects of biodiversity conservation- namely the importance of irreplaceability and complementarity. We know from multiple spatial prioritization studies on all kinds of taxa at different scales, that solutions that consider complementarity identify very efficient gains for protecting the pool of biodiversity- much more so than ranking places by species richness. For example, Madagascar

is usually very important contribution to the global species pool because of its high levels of endemism, but appears much less important here. Perhaps I've been looking at this map for too long, but I wonder if areas that are naturally sparse in forested vegetation are a bit inflated here (northern boreal zone in Canada, the mallee in South Australia/W. Victoria in Australia). I guess I'm struggling to see how thousands of species could be protected in northern Quebec (with very low endemism and species richness), and how
it could (or whether it should) rank more highly than protection of species in Madagascar.

RESPONSE: Thank you for this comment. You have raised a really important issue about complementarity. We have extended our model to account for overlap in sets of species across ecoregions. Responding to this comment was probably the most difficult of all of the comments raised by the referees but we agree that it is of fundamental importance. To address complementarity, we developed a probabilistic model that maximizes the expected number of unique species protected given protection levels and the sets of ecoregions occupied by different 'profiles' of species. Details of this extension are in the Methods ("Extensions to the base model"). We discuss the model in the main text in lines 88-92:

"We account for complementarity by developing a measure of species overlap among ecoregions. We then develop a probabilistic model that uses overlap to estimate the expected number of species that are protected with additional forest protection in an ecoregion that are not already protected elsewhere."

We show the effects of incorporating complementarity on optimal spending in a modified Figure 3 (panels A and D). We discuss these results in the main text in lines 149-155:

"Accounting for patterns of species overlap means there is greater benefit to conserving one representative ecoregion within a biome more completely while protecting less in nearby ecoregions (Fig. 3A). For example, in the Congo Basin, protected area increases dramatically in the Central Congolian Lowland Forests ecoregion, while protection declines in nearby ecoregions. Ecoregions with many endemic species and no closely related ecoregions (e.g., Hawaii) tend to have higher priority when complementarity is considered."

As for the specific locations you mention, accounting for overlap does increase area protected in one ecoregion within Madagascar, but there are declines in other ecoregions in Madagascar, in line with the discussion in the text. The effects of accounting for species overlap on forest area protected in Canada are mixed.

While, I agree that the financial realities are important in conservation decisions, I wonder about the use of a GDP-averaged metric to represent the cost of land. Some ecoregions might have inexpensive land from a global perspective, but that doesn't reflect the costs to the local people relying on those areas for persistence.

RESPONSE: Thank you for raising this important point. In our discussion paragraph detailing the limitations of our cost measure, we have acknowledged this aspect (lines 232-235):

“Issues with gridded GDP include being an incorrect signal of the value of forest land if the value of economic activity is not primarily related to forests or to land use, significant subsistence or cultural use values not reflected in GDP, and not incorporating maintenance and enforcement costs.”

An increase in such costs will have predictable consequences, tending to move forest protection away from those areas with increased costs. However, we are unaware of a global dataset that would allow us to adequately capture these valid concerns.

It’s pointed out that maintenance costs of conservation are not considered, but I also wondered about conservation effectiveness.

RESPONSE: Thank you for raising this issue. The extension we introduce to allow for species complementarity (see response above) is well-suited to accounting for variation in conservation effectiveness. Specifically, we recast additional protected area as increasing the probability of protection for a species in an ecoregion. That probabilistic framing allows for conservation effectiveness to vary by ecoregion, species, or even specific ecoregion-species combinations. The probabilistic approach could be used to study the effects of imperfect enforcement or maintenance of protected areas. That said, we are unaware of any reliable global datasets that we could use to parameterize those relationships to account for varying conservation effectiveness.

Minor Comments

Table S5 would be better as a data file, and it would also be nice to see the code for this paper.

RESPONSE: If the paper is accepted, we will include the revised Table S5 (now Table S3) as a data file and the code will be posted to a public repository. We are happy to share the code with referees of the paper now if they’d like.

It’s slightly confusing from the writing (e.g. abstract L12-13) whether the sequence of initially targeting ecoregions with little forest remaining+big gains in that remaining forest, then targeting other regions is intentionally done or whether it simply arises from the algorithm as the ‘final return-on-investment’ is always higher in these cases with a truncated trajectory (Fig. S1B). Obviously, there is a space limitation, but given this is a main result, perhaps a couple more sentences of explanation in the main text would help. Figure S1a is quite helpful conceptually, and I would think about adding it to the main text as well.

RESPONSE: We have clarified the language to say the optimization leads to spending initially on priority ecoregions (lines 77-83):

“…in some “priority ecoregions” with little remaining unprotected forest, relatively high species richness and low cost, the conservation ROI remains higher than in other ecoregions even after all remaining forest has been protected. In this case, it is optimal to devote the entire budget to protecting these priority ecoregions until no unprotected forest remains before beginning to protect other ecoregions. This initial allocation of all budget to a subset of ecoregions is an outcome of optimization (see SI Fig. S1 for intuition)."

We have also explained the logic of the optimal solution in the Methods section (lines 341-361) and in the SI in Section S1. We thought about including Figure S1 in the main text but opted not to for space reasons; however, we agree that it can be helpful for readers to think through this figure. If it is felt that it should be included in the main text we will work on figuring out how to include it given space constraints.

L149-159 I think this surrogacy analysis needs more explanation. What exactly is the conservation value here called 'optimal species protection'? If the number of protected species is the metric, then ranking by species richness should provide the optimal solution, and there should be some loss in that if costs/threats are taken into account. Especially given the uncertainty in the cost of conservation, shouldn't the current results be compared to an iterative ranking of ecoregions that maximizes biodiversity protection irrespective of costs? There should be trade-offs in either direction, how much extra would we have to spend to optimally protect species, and how many species do we lose if we budget.

RESPONSE: The objective function in our analysis is to maximize the number of protected species. By construction, the optimal solution is that which solves the budget-constrained species conservation problem as formulated in the Methods section. The number of species conserved using that approach is the optimal number of species which can be protected subject to budget constraints. We have added the "budget-constrained" modifier in front of the first instance of optimal species protection in that paragraph to reiterate that our problem is a constrained optimization problem.

As for your more specific comments about costs, we offer three responses. We see three possible different points you could be making and want to respond to each.

First, if you are suggesting that, with a finite budget, we would achieve higher levels of species protection using a prioritization scheme that ignores costs, we respectfully disagree. What is now Figure 2, panel B compares an approach which spends the marginal dollar in the area with the highest species richness (lime green, dot-dash line) to the optimal solution, which does account for costs (purple, solid line). The species-richness-only heuristic only achieves 22% of the optimal solution. Focusing on marginal gains in protected species richness (but still ignoring costs) does better, but still poorly compared to the optimal solution (47% of species gains). The cost of land is quite important in budget-constrained problems. We are not the first to make this point (e.g., Ando et al. 1998, Naidoo et al. 2006).

We believe you are also pointing out that having a budget constraint does reduce the number of species that can be protected as compared to a case with no budget constraint. We of course agree, and Figure S3.1 traces out the number of species that can be protected at different budget levels.

We also note you raise the issue of protection under uncertain costs. We agree this is an interesting avenue for research, but one we opt not to pursue in this paper. As detailed in responses to other reviewers raising the issues of uncertainty and robust optimization, properly accounting for uncertainty in model parameters dramatically increases the complexity of the problem. The optimal solution would need to account for the potential for both passive and active learning. While the role of costs would certainly change in such an approach, failing to consider *beliefs* about costs would still render a heuristic focused solely on species richness suboptimal. Information about costs in such a setting would be quite valuable, as conservation would still be bound by budget constraints, even if we are uncertain as to the costs.

L 214 – why is it either or ? why not look at endemics separately in the analysis?

RESPONSE: Thank you for raising this concern. This comment, your comment discussed above, and those from other reviewers led us to extend our model to account for complementarity (in the sense that a species can occur in multiple locations, not complementarity across species in a given ecoregion). As explained in the revised methods ("Extensions to the base model"), we now account for complementarity in a probabilistic framework. In that problem formulation, our goal is to maximize the expected number of additional species protected by expansion of protected forest area, accounting for the presence of some species in multiple ecoregions. In our complementarity framework, endemics will certainly have a different impact on conservation choices than other species, but their conservation is not the sole objective and their uniqueness fits as a special case within our broader framework. While we note that our approach could be applied equally to an objective of conserving endemic species alone, we find overall species conservation to be a more relevant objective given the public discourse around biodiversity conservation. For example, neither the goals nor targets in the Aichi Biodiversity Targets treat endemics separately.

We also note that the spatial protection patterns we observe when treating all species *as if* they were endemic align well with maps of plant endemism in prior work (Sandel et al. 2019).

I recommend considering some of the other global priorities/endemism papers and citing some of them, and perhaps comparing results in terms of priority areas and gains in efficiency:

Pouzols et at. 2014. Nature 18:383 (they also use ecoregions)

Pollock et al. 2017. Nature 546:141

Hobohm et al. 2019. Perspectives in Ecology and Conservation (endemic-area relationships of plants, which could be used here instead of species-area)

Sandel et al 2019 GEB 29:4

RESPONSE: Thank you for this suggestion. We have included a comparison to Pouzols et al. (2014) and Pollock et al. (2017) in Discussion and in the SI (new section S5.2). In the Discussion, we include a paragraph summarizing the comparison (lines 175-182):

"Other recent analyses of global conservation priorities using data on terrestrial vertebrates and ignoring cost and dynamics highlight several priority areas in common with our approach (e.g., Melanesia) but also significant differences. For example, our analysis prioritizes conserving forests in the Anatolian Peninsula, largely because these are small forest ecoregions with many plant species that can be protected at modest total cost. Conversely, we do not prioritize southeastern Brazil because of high land costs, which are likely due to proximity to large urban areas (see SI Section S5.2)[26,27]. Protecting animal species rather than plants may, of course, lead to different solutions."

We go into a bit more depth about these comparisons in the new SI Section S5.2:

<u>"S5.2 Comparison to other spatial prioritizations</u>

Our analysis is not the first to suggest a spatial conservation prioritization at a global scale. Recent high-profile examples include Pouzols et at. (2014)[20] and Pollock et al. (2017)[21], which offer interesting contributions in the form of static future land constraints, limitations of decentralized planning, and consideration of other biodiversity objectives beyond richness-based measures (e.g., functional or phylogenetic diversity). Those studies adopt a pixel-based prioritization approach using Zonation software, for which the additive benefit function objective most closely maps onto our objective.

While both our approach and these prior approaches identify some common regions as high priority (e.g., Melanesia), there are also clear differences. For example, our approach prioritizes areas in the Anatolian Peninsula far more than other approaches, while placing less emphasis on others (e.g., southeastern Brazil).

While a number of differences in modeling choices could explain these discrepancies, we note that Pouzols et al. (2014) and Pollock et al. (2017) do not account for costs in their prioritization. In fact, priority regions in those analyses align well with maps of tree endemism (Sandel et al. 2020[22]), even though the prioritization studies focus on animals. Regions where our approach places less emphasis than these other prioritization efforts have higher land costs as proxied by the Kummu et al (2018) data, often due to the presence of large urban areas (e.g., Rio de Janeiro, São Paolo, and Belo Horizonte in southeastern Brazil)."

We refrain from a quantitative comparison of efficiency as we have different objectives (plant species richness vs. those considered in the other studies, which focus on different taxa, and/or different biodiversity measures) and different measures of costs (monetary vs land area). Instead, we prefer to focus on the qualitative features of the spatial prioritization.

While we find the Hobohm et al. (2019) study an interesting contribution, we have opted not to refer to it directly in the paper. Given the excellent comments from all reviewers, we have extended our model to allow for complementarity across ecoregions, and as such do not wish to place as much emphasis on endemism.

Reviewer Reports on the First Revision:

Referees' comments:

Referee #1 (Remarks to the Author):

I want to thank the authors for carefully considering my comments and making several changes in the manuscript to address the comments. There are several improvements in this version of the paper. However I am still on the fence regarding the usefulness of protected species a the end of the planning horizon as the goal and I have some question about the answers to some other comments I raised. I discuss those comments below (using the original numbering I had used in my previous comments).

1. I find unsatisfactory that the authors say in their response that in the long run there is no difference between species protected and species remaining. It's a trivial statement that after 700 years at the deforestation rates used all forest will be gone or protected and in that case the two objective functions (species remaining, that I propose, and species protected, used by the authors) are the same. But to assume that deforestation rates will continue as they are for 700 years is obviously incorrect and in addition irrelevant for planning. So, to then state, that solutions for >50 years planning horizon are similar to the 700 years using their objective function, is in my opinion misleading. The question is how different is a solution for 10-50 years that maximize species remaining at the end of the planning period from a solution that maximizes species protected at the end of the planning period. And this fundamental issue is not addressed by the authors, other to say that "it is not well-behaved" mathematically and that there may be many "optimal" solutions for this other problem. I am wondering if this is, in the end, the main result of this analysis: that the policy relevant problem (species persisting at the end of the planning horizon) does not have a unique solution and it is not best addressed by such optimisation approach.

3. I appreciate that the authors now use SSP3 forest projections for an extra analysis in the SI. This scenario, as the authors state, has some of the highest forest loss among SSPs. A more contrasting case to their analysis would be SSP1 where there is indeed a decrease in the rate of deforestation over time. In any case, I do take the point that for this analysis, the GDP signal on land cost may dominate over changes in deforestation rates (and this could also be the case in SSP1 with higher GDP growth).

5. I think the authors are still not sufficiently critical of a global planning approach. Instead, they revise the text to imply that a national base approach would be non-optimal. There would be so many socio-political hurdles for having a central planning authority buying land for conservation globally that one could argue that from a broader policy view, that is what is not optimal (one just needs to write a different goal function that would include social conflict created by inequity in land access and price increases driven by such land acquisition). A deeper reflection on the limitations of this approach would really be beneficial to the readers of the paper. Otherwise the paper is vulnerable to recent criticisms raised about this type of global analysis.

Referee #2 (Remarks to the Author):

The authors have provided detailed responses to reviewer comments, and have made substantial revisions to address those concerns. The authors also now better acknowledge the assumptions and limitations of their analysis, which is essential.

Further concerns about the revised analysis:

1. L84 & 147. Why isn't the model that includes complementarity, reforestation and land price increases the primary model that is presented? Why give precedence to the more basic model that lacks these key considerations?

2. L67. "protected via purchase" The premise that purchasing land results in protection was a concern I raised previously. The authors argue that data availability prevents consideration of ongoing management and enforcement costs. However, doesn't this model just assume that protection costs are proportional to gridded GDP? If so, you do not need to specify 'purchasing land' as the mechanism that achieves protection. You could instead just state the assumption that protection costs are proportional to gridded GDP, which seems reasonable to me. This would also resolve many other problematic issues associated with purchasing land as the mechanism of protection: conservation colonialism, leakage, creating conservation solutions that have local community support, etc. The mention of land purchasing is unnecessary as you can adopt more general language without impact on the analysis.

3. Part of the novelty of this work is that it solves a scheduling problem: not only where to invest but when to invest (L2). Yet little attention is given to the sequencing of investment in the rest of manuscript other than reporting results in the first 10 years vs 50 years and noting the set of 68 priority ecoregions (L168). Yet a clear implication of the ordering is that new approaches to globally coordinated conservation are required to achieve that ordering. Much of the discussion focuses on technical issues, with only a few cursory remarks in the final paragraph about the policy implications of this work. By failing to more fully explore the implications of the scheduling aspects of this work the authors have missed an opportunity to highlight its novelty. And what is the cost of not achieving that ordering?

Minor comments

L63 "these data"?

L101 "additional" to what? It is not clear what the reference condition is that this statement implies. Is there some business-as-usual reference model or is the reference doing nothing? (This issue also applies to "gains" L102, L106 - a reference condition is implied).

L138 "the limited role of threat". This phrase is not clear in the context of this sentence.

L175 Consider adding citations to clarify which other analyses you are referring to here. And regarding L182, the difference may also arise from the use species-specific data, not species richness.

L241 Improve the effectiveness of conservation relative to what? ROI is already the standard that is considered best practice for conservation planning. L241 seems to imply that is not the case.

Referee #3 (Remarks to the Author):

Review of revised MS:
"Where to protect first in global forest conservation" by Ian Luby, Steve J. Miller, Stephen Polasky
Manuscript number 2021-02-01982A-Z

General Comments
The paper highlights the importance of trying to focus conservation resources upon those priority areas where they will be most efficiently applied, generating the greatest impact. While the base case is rather abstract the subsequent relaxing of assumptions makes this a very useful and interesting analysis. While I would like the authors to consider the specific comments made below and attend to them as they see fit, I would then support publication.
The most important change I feel is needed might well be considered one of the most trivial: the writing style needs improving. The paper is very clear in some places but really rather unclear in others. There are a number of places, especially early on, where the reader finds it difficult to follow what is actually going on – and yet then at other points it is crystal clear. If it could all be brought up to the same standard that would be greatly appreciated. This may seem trivial – but I value clarify greatly.
There are several places, highlighted in my specific comments, where the reader would greatly benefit from a more self-critical presentation, explaining the weaknesses in the assumptions would actually make this a stronger paper. For example, the approach taken effectively treats each ecoregion as internally homogeneous at its mean values. This is obviously quite a strong assumption within a given ecoregion – but I feel its defensible given that this is a global analysis. I doubt the authors themselves believe that the governments of the world, if briefly united, should follow the results of this study to the letter. It must be that, even within many of the highest priority ecoregions there are areas that would be highly inefficient for conservation. Making this clear (and with it highlighting the need for more micro-scale extension of these techniques at an ecoregion scale) would help the reader understand what is being assumed here. This won't at all devalue the study; on the contrary it will strengthen it.

A suggestion for future research
The above plus the specific comments below are all I would ask for in terms of changes to the paper. However I would like to suggest an extension for your future work. The focus on biodiversity would I think be more than enough for most referees and readers. However, my personal view is that it is in principle just as susceptible to the criticism of single focus decision making as could be applied to disastrous policies such as the Common Agricultural Policy (certainly for the first three decades of its operation). The real world is not single focus. It is the meeting of two hugely complex systems, one

environmental the other human. Each system is linked to the other in numerous, generally unplanned ways. This means that a single focus policy tends to have massive unforeseen consequences. The authors might argue that a focus on biodiversity would generate unforeseen positive benefits and certainly some positive externalities are to be expected. But there are also very likely to be negative consequences as well. The 'Food Security' wing of the agricultural lobby deliberately overplay their hand at every opportunity – but as with most long standing critiques there is a grain of truth there.
Now the present paper does a lot better than most. The inclusion of the cost variable helps to bring agricultural externalities into the prioritisation (even if it would be a crude approximation in a cost-benefit assessment; but this is a prioritisation telling us which is best, not telling us if any of these projects are cost-beneficial). However, there are many other impacts that are excluded. The one that might be considered further concerns greenhouse gas emissions and storage. It seems to me that this might be included in a further extension of this analysis without too much difficulty. One might argue that the change in net greenhouse gases was linearly related to the change (positive or negative) on the area of woodland conserved or lost. One could use a simple average quantity for this or use values that vary by biomes. Perhaps more attractively the authors could use a standard social cost of carbon value to convert this into an economic value and roll this into the cost function. This would still preserve the simplicity of a prioritisation analysis based on the cost-effectiveness of biodiversity conservation and avoiding the problems of getting into a cost-benefit frame.
The introduction of a linear net greenhouse gas removal value function which will have constant marginal values will contrast nicely with the declining marginal value curve assumed for biodiversity. This will illustrate a further very important issue regarding optimisation (one which I feel justifies the separate second paper); namely that changing the scope of impacts considered in an analysis can substantially change its outcomes. It is perfectly plausible that this extension might produce a considerably different outcome. Taking the two papers together could produce a very nice demonstration of optimisation principles: (i) The present paper: spatial prioritisation radically improves the efficiency of available resources (ii) A follow up paper: improving the incorporation of the complexities of the real world within analysis can result in significant changes to that prioritisation and further improve resource efficiency.
I don't feel this is a necessary requirement for the present paper but I hope the authors might find it an interesting idea for their future research.

Specific Comments
PLEASE SEE UPLOADED REVIEWER ATTACHMENT FOR MY SPECIFIC COMMENTS

Referee #4 (Remarks to the Author):

Your efforts in sorting out priorities for conservation of tropical forests are original and well worth pursuing. The quality of data a method of analysis are complete and compelling. The conclusions seem to me to be robust and very useful in this point of overall analysis of evidence in this situation. I have no further references to suggest and no negative comments on clarity.

Referee #5 (Remarks to the Author):

I have re-read the manuscript and all responses to reviewer comments. The authors did a very thorough job of responding to reviewer comments, and this version of the paper is a more robust contribution. I especially appreciate the authors considering my comments to add complementarity, including new analyses and results. I believe this will be a strong contribution to the global conservation literature, and I have no further comments.
Laura Pollock

Review of revised MS:

"Where to protect first in global forest conservation" by Ian Luby, Steve J. Miller, Stephen Polasky

Manuscript number 2021-02-01982A-Z

General Comments

The paper highlights the importance of trying to focus conservation resources upon those priority areas where they will be most efficiently applied, generating the greatest impact. While the base case is rather abstract the subsequent relaxing of assumptions makes this a very useful and interesting analysis. While I would like the authors to consider the specific comments made below and attend to them as they see fit, I would then support publication.

The most important change I feel is needed might well be considered one of the most trivial: the writing style needs improving. The paper is very clear in some places but really rather unclear in others. There are a number of places, especially early on, where the reader finds it difficult to follow what is actually going on – and yet then at other points it is crystal clear. If it could all be brought up to the same standard that would be greatly appreciated. This may seem trivial – but I value clarify greatly.

There are several places, highlighted in my specific comments, where the reader would greatly benefit from a more self-critical presentation, explaining the weaknesses in the assumptions would actually make this a stronger paper. For example, the approach taken effectively treats each ecoregion as internally homogeneous at its mean values. This is obviously quite a strong assumption within a given ecoregion – but I feel its defensible given that this is a global analysis. I doubt the authors themselves believe that the governments of the world, if briefly united, should follow the results of this study to the letter. It must be that, even within many of the highest priority ecoregions there are areas that would be highly inefficient for conservation. Making this clear (and with it highlighting the need for more micro-scale extension of these techniques at an ecoregion scale) would help the reader understand what is being assumed here. This won't at all devalue the study; on the contrary it will strengthen it.

A suggestion for future research

The above plus the specific comments below are all I would ask for in terms of changes to the paper. However I would like to suggest an extension for your future work. The focus on biodiversity would I think be more than enough for most referees and readers. However, my personal view is that it is in principle just as susceptible to the criticism of single focus decision making as could be applied to disastrous policies such as the Common Agricultural Policy (certainly for the first three decades of its operation). The real world is not single focus. It is the meeting of two hugely complex systems, one environmental the other human. Each system is linked to the other in numerous, generally unplanned ways. This means that a single focus policy tends to have massive unforeseen consequences. The authors might argue that a focus on biodiversity would generate unforeseen positive benefits and certainly some positive externalities are to be expected. But there are also very likely to be negative consequences as well. The 'Food Security' wing of the agricultural lobby deliberately overplay their hand at every opportunity – but as with most long standing critiques there is a grain of truth there.

Now the present paper does a lot better than most. The inclusion of the cost variable helps to bring agricultural externalities into the prioritisation (even if it would be a crude approximation in a cost-benefit assessment; but this is a prioritisation telling us which is best, not telling us if any of these projects are cost-beneficial). However, there are many other impacts that are excluded. The one that might be considered further concerns greenhouse gas emissions and storage. It seems to me that this might be included in a further extension of this analysis without too much difficulty. One might argue that the change in net greenhouse gases was linearly related to the change (positive or negative) on the area of woodland conserved or lost. One could use a simple average quantity for this or use values that vary by biomes. Perhaps more attractively the authors could use a standard social cost of carbon value to convert this into an economic value and roll this into the cost function. This would still preserve the simplicity of a prioritisation analysis based on the cost-effectiveness of biodiversity conservation and avoiding the problems of getting into a cost-benefit frame.

The introduction of a linear net greenhouse gas removal value function which will have constant marginal values will contrast nicely with the declining marginal value curve assumed for biodiversity. This will illustrate a further very important issue regarding optimisation (one which I feel justifies the separate second paper); namely that changing the scope of impacts considered in an analysis can substantially change its outcomes. It is perfectly plausible that this extension might produce a considerably different outcome. Taking the two papers together could produce a very nice demonstration of optimisation principles: (i) The present paper: spatial prioritisation radically improves the efficiency of available resources (ii) A follow up paper: improving the incorporation of the complexities of the real world within analysis can result in significant changes to that prioritisation and further improve resource efficiency.

I don't feel this is a necessary requirement for the present paper but I hope the authors might find it an interesting idea for their future research.

Specific Comments

Line	Text	Review comment
3	Optimization	While it not needed in the summary paragraph, it is worth making clear to the reader that there is not just one approach to optimisation. Approaches such as greedy algorithms and

		combinatorial optimisation can give significantly different answers, especially when applied to a large set of options such as global forestry protection. The approach adopted needs clarification and justification.
5-7	The optimal conservation strategy accounting for species richness, costs of conservation, existing degree of protection, and threat of habitat loss	There is a need to clarify to the reader that there is no universal optima; rather the optimal solution changes with every reformulation of the objective – and there are multiple defensible objectives. If we argue that the use of land cost proxies the forgone benefits of market production such as agricultural output then perhaps the most important omission is in terms of carbon storage (mentioned further below). Including this would considerably change optimal outcomes.
4 vs. 8 vs. 10	global priorities for expanding protected areas vs. priority ecoregions where further forest loss leads to large reductions in species conservation vs. conservation efforts	You switch from talking about optimising for expanding protection to discussing further forest loss and then switch to conservation efforts. These are three different types of action within 6 lines presented as if they are the same thing. But these are not the same thing and you would undertake these different objectives in alternative locations. This is rather confusing.
12	A strategy using conservation return on-investment while ignoring threat	This seems like a fourth action – which is now really confusing as is the concept of "conservation return on-investment while ignoring threat" which is not something I intuitively understand. The term 'threat' is also not defined here.
14	equal fractional protection targets	I don't know what this means – and given I have worked in this field for a while I think this might be a common reaction to this term
11-14	A strategy using conservation return on-investment while ignoring threat is close to optimal, while focusing solely on species rich areas, low-cost areas, highly threatened areas, or equal fractional protection targets achieves only 5-47% of potential conservation gains.	I get the general gist of this but overall I don't really know what this is trying to say. Even terms such as "potential conservation gains" are not self-evident. Can the summary paragraph be rewritten so that, as is required, it is understood by the general Nature reader.

14 - 17	Including the potential for reforestation further emphasizes the importance of priority ecoregions, while including land costs that rise with land scarcity works in the opposite direction.	Taken at face value this sentence does not make sense. Read literally it says that including the cost of land reduces the importance of priority ecoregions.
1-18	Summary paragraph.	Overall the summary paragraph needs a thorough rewriting to make it generally accessible, clear and unambiguous.
24	virtually all major habitats have seen declines	Technically this cannot be true. What has happened is that natural habitats have been progressively replaced by developed habitats.
27-29	Since 1990, 178 million hectares of forest have been lost, representing 4.2% of forests that remained in 1990.	Minor: Rewrite to only mention 1990 once.
51-52	We used data on 458 forest ecoregions based on The Nature Conservancy's (TNC) global terrestrial ecoregions	Please clarify the extent of the analysis – are you considering all terrestrial area? Given the uncertainty in the summary paragraph its not clear to me if you are considering all forests, all areas that are or could be/where forests or some other area.
56-57	and the number of plant species (TNC). We used plant species data because plants are an important component of overall biodiversity	I am uncertain how this measure is working in practice. Is it as follows: Suppose we have three areas, A, B and C with species as follow (trees used for illustration): Area A contains oak, ash and willow Area B contains oak, ash and willow Area C contains elm and birch So using the scoring procedure described in the paper the biodiversity measures are as follows: A = 3 B = 3 C = 2 Therefore it seems to me that the optimisation process will select one of A or B to start off with. Lets suppose it picks A Then it will select B because 3>2 Finally it selects C

		However if resource constraints force us to only choose to protect 2 sites then only A and B are chosen. This satisfies the optimisation rule but leads to a poorer outcome because these two sites only protect 3 species rather than the 5 that would be protected is A and C were chosen. Have I misunderstood what is happening here? If so then this needs clarifying. If not then the present approach needs justification. Later on (line 85) you state that the approach is "implicitly assuming that all species are endemic". I think this is perhaps how you avoid this problem – effectively by assuming that the commonality of species in Sites A and B (above) can never happen. Is this correct? If so then perhaps you should discuss the challenges that arise when that endemism assumption does not hold. I'm not an ecologist so I can't judge how realistic the endemism assumption might be in a world split into just 458 forested Ecoregions – maybe that's perfectly defensible.
65-66	The number of species protected in each ecoregion is a power function of the forest area protected in that ecoregion ("species-protected-area curve").	This is defensible as far as it goes but I would like to see more self-critique. If you are optimising biodiversity and defining that in terms of the number of species conserved then you are ignoring abundance. Does this means that a state in which there are just a small number of each of X species will be preferred over a great abundance of X-1 species? If so then I think this needs clarifying as it would be debateable whether such an objective is what we want to optimise.

		Also, if we are optimising something that follows a power curve how does this relate to other benefits from the land which follow different relationships? For example, agricultural produce might be best approximated by a linear function. It seems plausible that this is captured in the linear cost function. But what about carbon storage? This might also be best approximated in a global analysis by a linear relationship with area – but its non-market value won't be captured by the GDP proxy. Does this matter? Perhaps not because the exercise is looking for a prioritisation i.e. it is answering a question about which order some pre-determined action should be taken in, rather than the question of whether action is actually justified. However I would really like to see some discussion of these issues.
73-74	Additional species protected per dollar spent	I refer back to my comment on lines 56-57; are these actually additional distinct species or more of the same but in different places? I really like this paper and don't want to tie it down to the point that its important message is lost but clarification of what is being done would help
88-89	We account for complementarity by developing a measure of species overlap among ecoregions.	Excellent – that addresses my concern above. I will leave those comments in just so you can see what readers reactions might be. You might think about heading those concerns off earlier just by flagging up that you test assumptions subsequently.
84-97		This paragraph did a lot to allay my concerns about the paper
102-106	The largest gains in species conservation come in Melanesia, Southeast Asia, parts of China, the Anatolian Peninsula, northern South America, and Central America (Fig. 1). In contrast, little forest is protected in many higher income regions with high land costs including the USA, Europe, Japan and Eastern	This makes a lot of sense as a result – but what is the relative strength of the cost variable in driving this result as opposed to perhaps higher biodiversity scores in tropical zones? A separate point is that many of these high income/high cost countries are

	China.	pushing afforestation targets for greenhouse gas removal. A common critique of this is that it would be more cost-effective for those countries to pay for afforestation to happen in lower income countries s well. Does your result suggest that there might be a biodiversity bonus to such outsourcing of afforestation? (it seems plausible to me).
110-111	In virtually all priority ecoregions (67 of 68), all land is either conserved or deforested within 10 years.	Does this imply that without action all of this land will be deforested in 10 years?
118-119	the rate of deforestation (threat)	This definition of threat could usefully have been clarified earlier.
129-130	the benefit of spending the 50-year budget up front. The gains from immediate access to the entire budget are modest, an additional 124 species (0.12%).	This is a surprising result – what is driving this? It seems at odds with the earlier result at lines 110-111 which suggested to me that timing really mattered. How can it be that, without conservation, all priority ecoregions that are not already protected disappear – yet making a 50 year budget all available in year 1 makes so little difference? Some discussion of how this apparent contrast arises would be helpful.
140	the heuristic targets equal fractional protection	I did not understand this phrase
147-160		Very nice sensitivity analyses yielding expected but pleasing results.
212-222		The discussion regarding the incorporation of climate change impacts differs from the point I raise regarding net greenhouse gas removal (above).
235-238	we have not incorporated heterogeneity within ecoregions that could allow for targeted conservation in areas with greater biological benefit, lower cost, or more immediate threats. As better data becomes available, the analysis can be redone to refine our results.	Good to see this. Perhaps an earlier note regarding this would be helpful in clarifying the approach taken, its strengths and limitations.
274-275	Figure 2. A) Relationship between the number of species protected in the first 10 years in each ecoregion (circle size)	Suggest this is rewritten as: Figure 2. A) Relationship between the number of species protected in each ecoregion (circle size) in the first 10 years

277-278	Threat is defined as the fraction of remaining unprotected forest that is deforested per year at current rates.	I would have appreciated seeing this definition earlier on in the paper.
296	Figure 3	The small size of these four global maps makes the information difficult to read. First impressions are that they are all very similar. Looking at them in the unconverted form the differences are much clearer – but that is in major part because the figures are simply much bigger there. This is a production issue (rather than one for the authors) but excessive reduction will render these figures unreadable.
43-45 vs. 452-475		On 43-45 the paper states “Here we apply a dynamic optimization approach to maximize the number of plant species conserved in the world’s forests” but 452-475 discusses “the abundance of profiles for mammals across the same ecoregions”. I found this confusing as I had assumed throughout my reading of the paper up to that point that the analysis was solely concerns with plants and that the concept of complementarity was also in terms of plants. I admit I am confused by this.

Author Rebuttals to First Revision:

Referees' comments:

Referee #1 (Remarks to the Author):

I want to thank the authors for carefully considering my comments and making several changes in the manuscript to address the comments. There are several improvements in this version of the paper. However I am still on the fence regarding the usefulness of protected species a the end of the planning horizon as the goal and I have some question about the answers to some other comments I raised. I discuss those comments below (using the original numbering I had used in my previous comments).

RESPONSE: We thank the reviewer for their comments on both the first and second round of review. We made extensive revisions to the paper in light of your initial set of comments and those of the other reviewers. We think these revisions have improved the paper. We have undertaken further revisions to the paper in response to the second round of comments, which we document below.

1. I find unsatisfactory that the authors say in their response that in the long run there is no difference between species protected and species remaining. It's a trivial statement that after 700 years at the deforestation rates used all forest will be gone or protected and in that case the two objective functions (species remaining, that I propose, and species protected, used by the authors) are the same. But to assume that deforestation rates will continue as they are for 700 years is obviously incorrect and in addition irrelevant for planning. So, to then state, that solutions for >50 years planning horizon are similar to the 700 years using their objective function, is in my opinion misleading. The question is how different is a solution for 10-50 years that maximize species remaining at the end of the planning period from a solution that maximizes species protected at the end of the planning period. And this fundamental issue is not addressed by the authors, other to
say that "it is not well-behaved" mathematically and that there may be many "optimal" solutions for this other problem. I am wondering if this is, in the end, the main result of this analysis: that the policy relevant problem (species persisting at the end of the planning horizon) does not have a unique solution and it is not best addressed by such optimisation approach.

RESPONSE: As we noted in our responses to the first round of comments, we thought long and hard about this issue. We remain convinced that our objective of maximizing protected species is superior to the alternative of maximizing (protected and unprotected) species at the end of the planning horizon. We apologize for not explaining our logic sufficiently in the prior round.

In a dynamic problem with a finite end date that occurs prior to the cessation of the dynamic processes being modeled (deforestation and protection), there is a question of how one should value conditions at the end date when there is additional value that will continue to accrue beyond the planning horizon. We chose to have a finite end date for the obvious reason that uncertainty in the far distant future makes the full planning horizon out to the end of the dynamic processes (i.e., when there is no longer any unprotected forest remaining, which is 725 years in the case without reforestation) an unrealistic exercise. Solving a dynamic optimization problem with a finite end date, however, doesn’t mean we should ignore what happens beyond the planning horizon. In the case of biodiversity conservation, we certainly care about what happens to species after the 50-year planning horizon is over. The typical approach in dynamic programming problems with a finite planning horizon and on-going processes is to consider the continuation problem after the planning horizon and show how the conditions at the end of the planning horizon are likely to influence value in the continuation problem. With our existing objective of maximizing the number of species in protected forests, we find that when we solve the model over longer time horizons (100 years, 200 years; see SI Section S3), species protection after 50 years for these longer planning horizon problems is similar to what we observe in the 50-year horizon solution, with protection in the first few years almost identical. It is in this spirit that we are convinced that the objective of maximizing the number of protected species at the end of the planning horizon is the most sensible approach.

Our point in analyzing a "sufficiently long planning horizon" to reach the logical endpoint (i.e., full planning horizon) is that there is no ambiguity in defining the objective because there is no difference in species in protected areas and species in total. We can use results for the full planning horizon to see how solutions in the first T-years of the full planning horizon compare with solutions of a T-year planning horizon. We disagree with the reviewer's statement that such a comparison is "misleading." However, we no longer mention optimizing to the full planning horizon in the main text. As noted above, we also found a close correspondence between the solution of our 50-year planning horizon with the first 50 years of a 100-year planning horizon and a 200-year planning horizon.

The alternative objective suggested by the reviewer is to maximize the number of species (protected and unprotected) at the end of the planning horizon. We agree with the reviewer that this alternative approach has some appeal because it is the existence of species, not whether they are formally protected or not, that should matter. Under this alternative approach, however, there is no gain in the objective from protecting remaining forest in any ecoregion unless it changes the amount of deforestation *within the planning horizon*. Since the finite end of the planning horizon is arbitrary, and since there is value to conservation efforts even beyond the planning horizon, not giving any value to additional protection in ecoregions with some amount of unprotected forest doesn't accord with intuition. In short, this objective fails to capture the value that forest protected during the planning horizon offers in the continuation problem, resulting in myopic decision-making in what is clearly a dynamic problem..

We have rewritten the paragraph in the discussion where we discuss the appropriateness of the objective trying to capture the logic explained above but doing so in a space-constrained manner:

"In our analysis, we assume the conservation objective is to maximize the number of plant species in protected forested areas globally at the end of the planning horizon. An alternative conservation objective is to maximize the total number of species in forested ecoregions whether or not they occur in protected areas. Both objectives are potential milestones for what we ultimately care about: existence of species far into the future. With our existing objective of maximizing the number of species in protected forests, we find a close correspondence between the solution of our 50-year planning horizon with the first 50 years of a 100-year planning horizon and a 200-year planning horizon (see SI Section S3). In contrast, maximizing species in forested ecoregions at the end of our planning horizon could lead us astray from our longer-term goal. There is no gain in this alternative objective to protecting forest in ecoregions with remaining unprotected forest at the end of the planning horizon, as protecting in those areas would not prevent deforestation. Many forest ecoregions therefore receive no investment, even though investment in the initial period is optimal when a longer time horizon is considered."

3. I appreciate that the authors now use SSP3 forest projections for an extra analysis in the SI. This scenario, as the authors state, has some of the highest forest loss among SSPs. A more contrasting case to their analysis would be SSP1 where there is indeed a decrease in the rate of deforestation over time. In any case, I do take the point that for this analysis, the GDP signal on land cost may dominate over changes in deforestation rates (and this could also be the case in SSP1 with higher GDP growth).

RESPONSE: We agree with the referee that using SSPs to benchmark future deforestation rates is a good idea, which is why we now include analysis with SSP3. We also agree in principle that showing outcomes under multiple SSPs can illustrate how different cost and threat levels affect results. However, using the most optimistic SSP (SSP1) as a benchmark is problematic. With only minor exceptions, SSP1 forecasts net increases in forest cover over the next 50 years (and no projections of gross deforestation rates). For the sake of biodiversity conservation, we hope this forecast comes true. For modeling purposes, however, this forecast makes the question of trying to prevent deforestation uninteresting.

In the spirit of the reviewer's comment, we have run two additional scenarios which examine reduced deforestation rates: (1) we use half of historical rates and (2) we use half of the SSP3 rates, keeping cost and budget increases at levels implied by SSP3. In both cases (figures below and added to the SI), we find overall spatial patterns of conservation to be similar to the results in the main text, especially for the areas in which all conservation spending occurs in the first year. Both figures contain the same information as those for robustness checks in the SI: the top panel depicts expected species conserved by ecoregion, with areas in which all conservation spending occurs in the first year highlighted in yellow, while the bottom panel depicts differences in expected species conserved by ecoregion compared to the model used in the main text.

Half of historical deforestation rates:

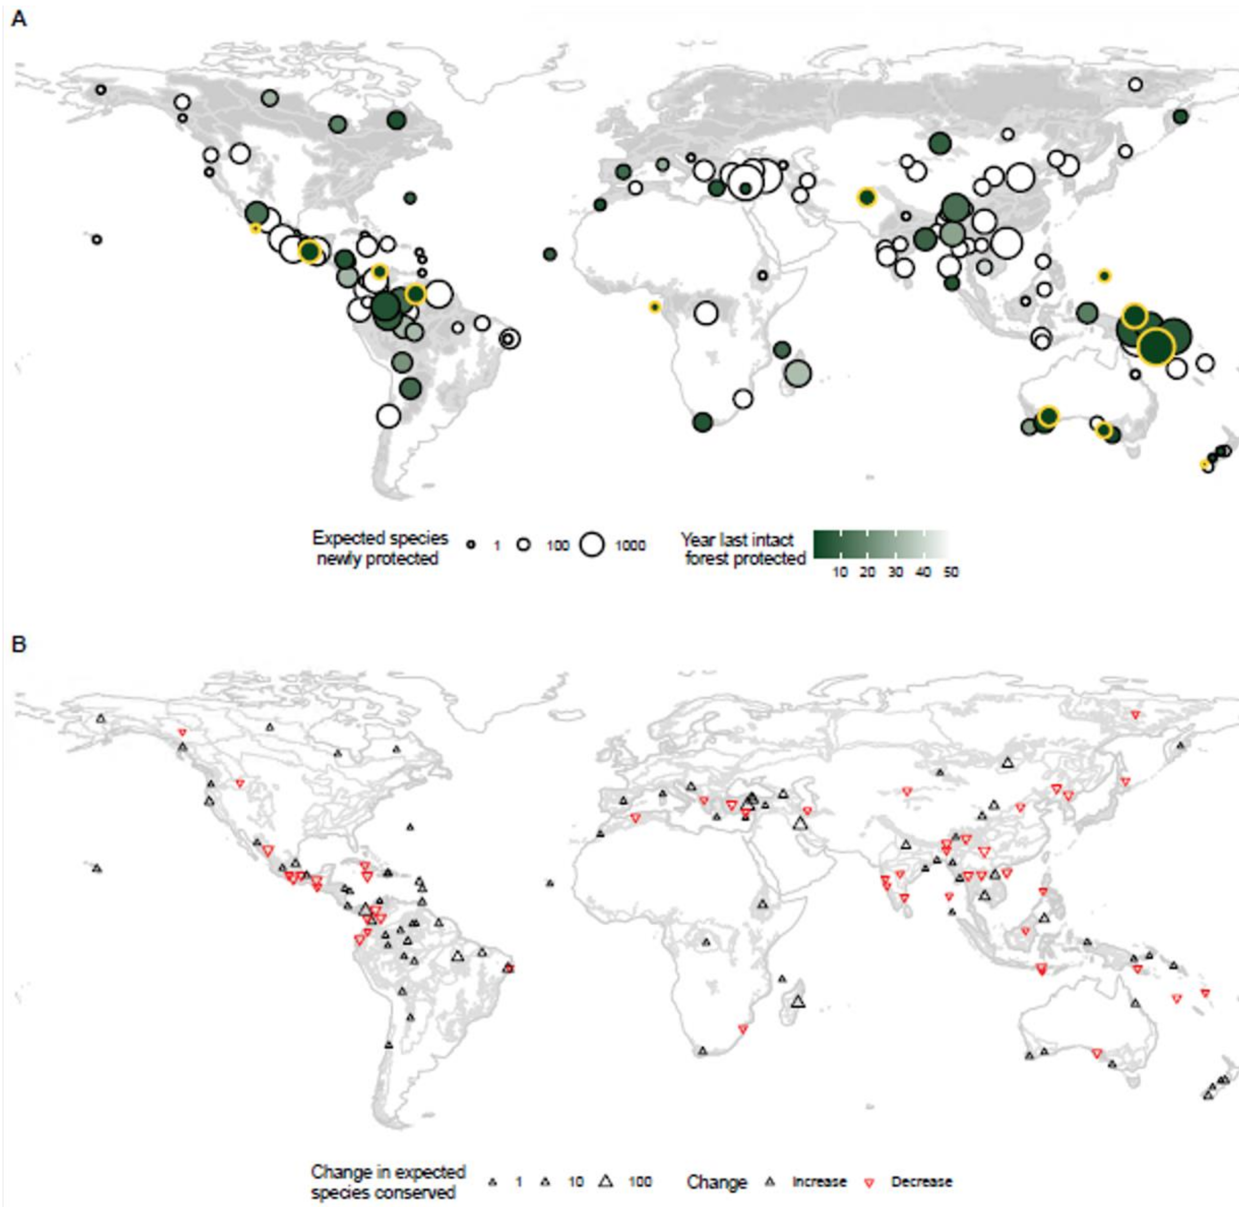

A
Expected species newly protected
1
100
1000
Year last intact forest protected
10
20
30
40
50
B
Change in expected species conserved
1
10
100
Change
Increase
Decrease

Half of SSP3 deforestation rates:

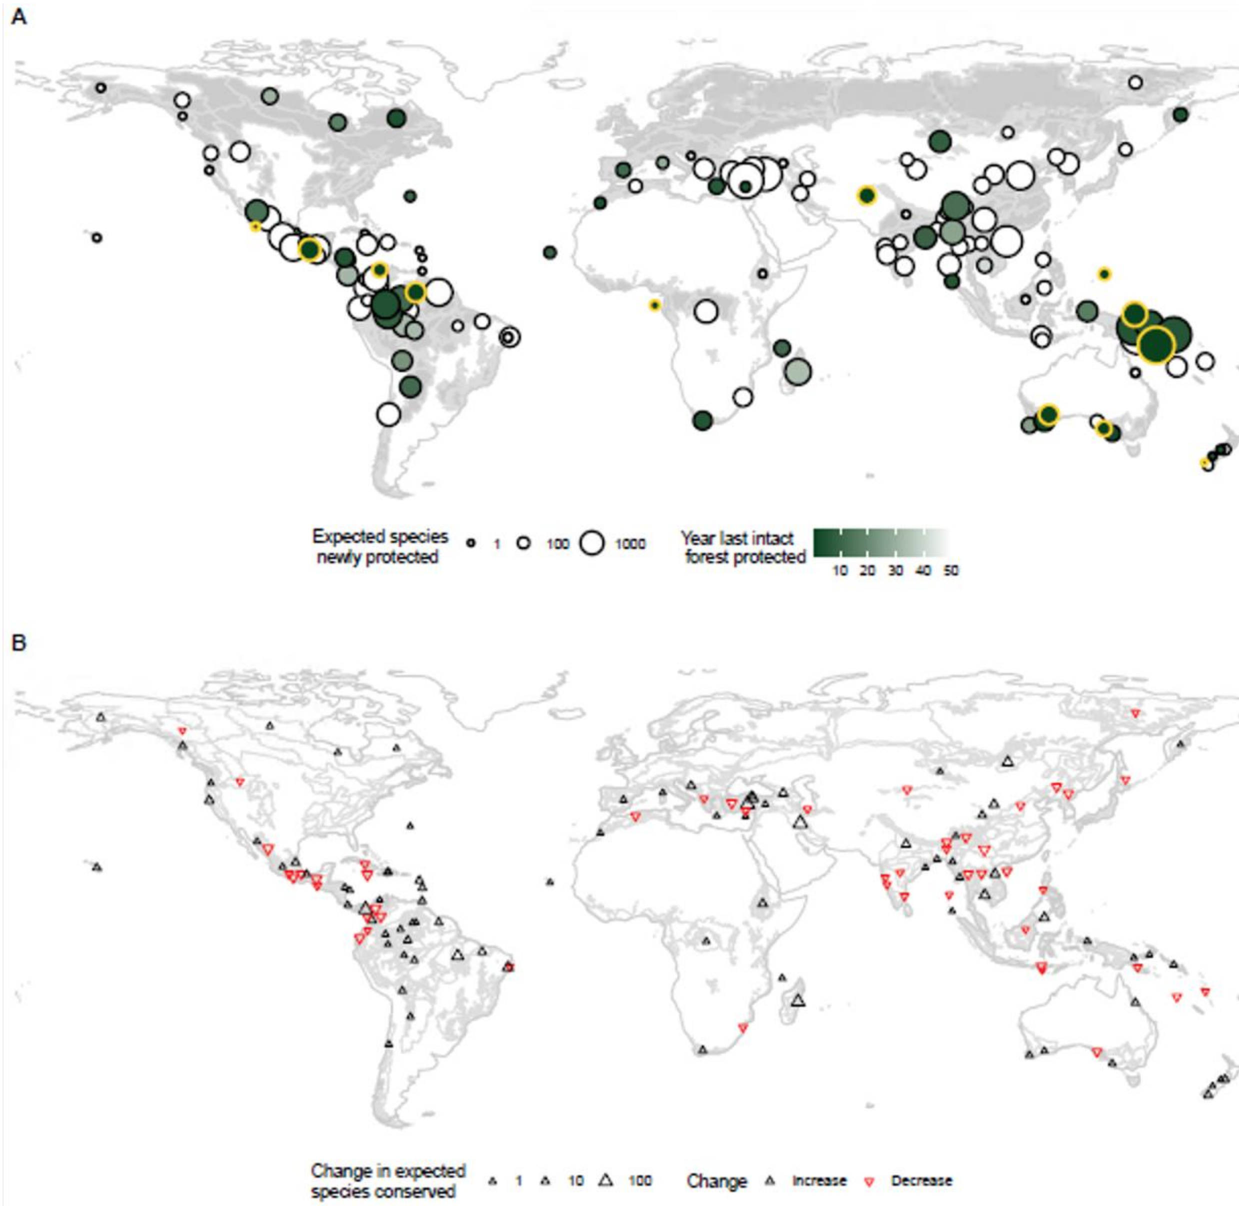

5. I think the authors are still not sufficiently critical of a global planning approach. Instead, they revise the text to imply that a national base approach would be non-optimal. There would be so many socio-political hurdles for having a central planning authority buying land for conservation globally that one could argue that from a broader policy view, that is what is not optimal (one just needs to write a different goal function that would include social conflict created by inequity in land access and price increases driven by such land acquisition). A deeper reflection on the limitations of this approach would really be beneficial to the readers of the paper. Otherwise the paper is vulnerable to recent criticisms raised about this type of global analysis.

RESPONSE: We agree with your comment. Our original response to your comment did not go far enough to address this concern. In the Discussion section we have rewritten the paragraph that talks about national or subnational approaches to conservation:

“Most decisions affecting conservation are made at national or subnational levels with little coordinated global conservation planning. The methods developed in this paper could be applied to conservation planning at the national or subnational scales with consequent improvements in such planning. We did not do so here because we lacked national-level budgets to inform the exercise and because we wanted to highlight what is possible with coordinated global planning. There is likely to be some reduction in effectiveness when conservation planning is done at a national rather than global scale because of possible failure to consider complementarity and the potential for unequal conservation ROI across countries[26]. The ability of international conservation organizations to move resources among countries, however, can help achieve some of the gains of coordinated global planning.”

Referee #2 (Remarks to the Author):

The authors have provided detailed responses to reviewer comments, and have made substantial revisions to address those concerns. The authors also now better acknowledge the assumptions and limitations of their analysis, which is essential.

RESPONSE: We thank the reviewer for their comments on both the first and second round of review. We made extensive revisions to the paper in light of your initial set of comments and those of the other reviewers, which resulted in an improved paper. In response to this round of comments, we have undertaken further revisions which we document here.

Further concerns about the revised analysis:

1. L84 & 147. Why isn't the model that includes complementarity, reforestation and land price increases the primary model that is presented? Why give precedence to the more basic model that lacks these key considerations?

RESPONSE: You make a good point, and we agree. This comment was the one that we spent the most time addressing in our revision and probably made the greatest improvements in exposition between the last draft and this draft. We have rewritten the paper so that the model that includes complementarity, land prices that rise with increased protection, and reforestation potential, is the primary model we discuss in the main text. We have removed discussion of our simpler model without these features from the main text. We hadn’t made the switch in the prior revision mostly out of inertia (we were still thinking of our original analysis as “the base case” and the additional features as “extensions”), which is not a good reason for keeping the old way of presenting the material. Focusing on the model with complementarity, rising land price with increased protection, and reforestation potential meant we had to extensively reorganize and rewrite the main text and the Methods section (changes too numerous to list out explicitly here). The Methods section in particular was thoroughly rewritten, resulting in a number of improvements in exposition.

2. L67. "protected via purchase" The premise that purchasing land results in protection was a concern I raised previously. The authors argue that data availability prevents consideration of ongoing management and enforcement costs. However, doesn't this model just assume that protection costs are proportional to gridded GDP? If so, you do not need to specify 'purchasing land' as the mechanism that achieves protection. You could instead just state the assumption that protection costs are proportional to gridded GDP, which seems reasonable to me. This would also resolve many other problematic issues associated with purchasing land as the mechanism of protection: conservation colonialism, leakage, creating conservation solutions

that have local community support, etc. The mention of land purchasing is unnecessary as you can adopt more general language without impact on the analysis.

RESPONSE: We agree. We have removed mention of land purchase from the paper and now discuss protection costs more generically as being proportional to gridded GDP.

3. Part of the novelty of this work is that it solves a scheduling problem: not only where to invest but when to invest (L2). Yet little attention is given to the sequencing of investment in the rest of manuscript other than reporting results in the first 10 years vs 50 years and noting the set of 68 priority ecoregions (L168). Yet a clear implication of the ordering is that new approaches to globally coordinated conservation are required to achieve that ordering. Much of the discussion focuses on technical issues, with only a few cursory remarks in the final paragraph about the policy implications of this work. By failing to more fully explore the implications of the scheduling aspects of this work the authors have missed an opportunity to highlight its novelty. And what is the cost of not achieving that ordering?

RESPONSE: Thank you for this comment. We agree that the prior revision missed the opportunity to highlight the novelty of the sequencing of investments, which comes from solving the dynamic optimization. We have made a number of changes to the paper to emphasize results on timing. In the abstract, we have rewritten the second and third sentence to emphasize timing:

"With limited resources and imminent threats, deciding when as well as where to conserve is a fundamental question. We use a dynamic optimization approach to identify sequencing of where to conserve plant species in 458 forested ecoregions globally over the next 50 years."

We have rewritten the second paragraph of the results section to give greater depth of discussion to timing:

"Conservation spending is highly concentrated in both space and time. In the first year, conservation investment targets just 18 ecoregions out of 458. In 13 of these ecoregions no unprotected forest remains at the end of the first year (highlighted in yellow in Fig. 1). These ecoregions have little remaining forest and high conservation ROI. Conservation spending expands to more ecoregions over the first 10 years but remains concentrated in only 10% of all ecoregions (46 out of 458). Shifting attention to temporal patterns, conservation gains in the first 10 years make up a large fraction of the species protected over the entire 50-year planning horizon (12,226 species, 51.6%). These broad patterns of optimal conservation are robust to a variety of alternative assumptions and data sources (SI Section S3). Where final species protection results differ, they do so in ways consistent with underlying data differences and the intuition provided above."

We have modified figure 1 to emphasize sequencing of investment. We use yellow highlights to show the 13 ecoregions where all unprotected forest is protected in the first year. We also provide shading to show when within the 50-year planning horizon that protection covers all remaining intact forest. Previously we had only shown whether this occurred within the first ten years or after.

We found that a heuristic approach that maximizes the initial conservation return-on-investment (ROI), defined as the initial benefit (species gain per hectare) divided by initial cost ($ per hectare), did almost as well as the dynamic optimal solution (Fig. 2B). Of course, we wouldn't have been able to show this result without being able to compare the heuristic approach with the optimal dynamic solution. This shows that a well-chosen (static) heuristic that bases where to invest currently on the basis of current ROI does almost as well as the optimal dynamic solution. We found it somewhat surprising that the rate of deforestation (threat) didn't play a larger role in determining the optimal sequencing of conservation investments. A heuristic prioritizing on the basis of threat alone performed the worst of all heuristics we analyzed.

Minor comments

L63 "these data"?

RESPONSE: In rewriting, we have removed this part of the sentence so it no longer refers to "these data." The new version of this sentence reads: "We applied a dynamic optimization approach to maximize the expected number of plant species protected in the 458 global forest ecoregions at the end of a 50-year planning horizon."

L101 "additional" to what? It is not clear what the reference condition is that this statement implies. Is there some business-as-usual reference model or is the reference doing nothing? (This issue also applies to "gains" L102, L106 - a reference condition is implied).

RESPONSE: this is additional to the case where there is no conservation spending. We have added the phrase "compared to no conservation spending" in the first sentence of the results.

L138 "the limited role of threat". This phrase is not clear in the context of this sentence.

RESPONSE: We have taken out "the limited role of threat" from this sentence. The main point of the sentence was about the role that well-designed heuristics could play, which is a point made previously by Wilson et al. (2006).

L175 Consider adding citations to clarify which other analyses you are referring to here. And regarding L182, the difference may also arise from the use species-specific data, not species richness.

RESPONSE: We have moved the references to Pouzols et al. (2014) and Pollock et al. (2017) up to the first sentence of this paragraph to clarify that these are the analyses to which we are comparing our results. We have also added several additional references that analyze global biodiversity priority setting. In the last sentence of this paragraph we now mention that having species-specific data, in addition to differences between plants and vertebrates, ignoring costs and threat, and inclusion of other objectives such as carbon, may cause differences in solutions.

L241 Improve the effectiveness of conservation relative to what? ROI is already the standard that is considered best practice for conservation planning. L241 seems to imply that is not the case.

RESPONSE: we agree with your point and have deleted this sentence.

Referee #3 (Remarks to the Author):

Review of revised MS:
"Where to protect first in global forest conservation" by Ian Luby, Steve J. Miller, Stephen Polasky
Manuscript number 2021-02-01982A-Z

General Comments
The paper highlights the importance of trying to focus conservation resources upon those priority areas where they will be most efficiently applied, generating the greatest impact. While the base case is rather abstract the subsequent relaxing of assumptions makes this a very useful and interesting analysis. While I would like the authors to consider the specific comments made below and attend to them as they see fit, I would then support publication.

RESPONSE: We thank the reviewer for their comments on both the first and second round of review. We made extensive revisions to the paper in light of your initial set of comments and those of the other reviewers, which resulted in an improved paper. In response to this round of comments, we have undertaken further revisions which we document here.

The most important change I feel is needed might well be considered one of the most trivial: the writing style needs improving. The paper is very clear in some places but really rather unclear in others. There are a number of places, especially early on, where the reader finds it difficult to follow what is actually going on – and yet then at other points it is crystal clear. If it could all be brought up to the same standard that would be greatly appreciated. This may seem trivial – but I value clarify greatly.

RESPONSE: We also greatly value clarity and apologize for not being sufficiently clear in parts of the revised draft. We did extensive rewriting that we hope makes the paper clear throughout. In response to your comments and those of reviewer 2, we refocused the paper on the model that includes complementarity, land prices that rise with increased protection, and reforestation potential. We have removed discussion of our simpler model without these features from the main text. This involved extensive rewriting of the methods and the results. We think this set of changes addresses many of the specific comments you raised (responses to those are below).

There are several places, highlighted in my specific comments, where the reader would greatly benefit from a more self-critical presentation, explaining the weaknesses in the assumptions would actually make this a stronger paper. For example, the approach taken effectively treats each ecoregion as internally homogeneous at its mean values. This is obviously quite a strong assumption within a given ecoregion – but I feel its defensible given that this is a global analysis. I doubt the authors themselves believe that the governments of the world, if briefly united, should follow the results of this study to the letter. It must be that, even within many of the highest priority ecoregions there are areas that would be highly inefficient for conservation. Making this clear (and with it highlighting the need for more micro-scale extension of these techniques at an ecoregion scale) would help the reader understand what is being assumed here. This won't at all devalue the study; on the contrary it will strengthen it.

RESPONSE: We recognize that global studies such as this necessarily simplify the problem; ours ignores important details such as heterogeneity within ecoregions. We discuss a number of important limitations in the Discussion section. We include a sentence about heterogeneity: "Incorporating heterogeneity within ecoregions would allow for targeted conservation in areas with greater biological benefit, lower cost, or more immediate threats." The Discussion section also talks about data limitations, climate change, and other considerations that could be included in the model. We also discuss how the objective function and use of specific data and assumptions affect the results, with some of this in the results section, the discussion section, and in the Supplementary Information.

A suggestion for future research
The above plus the specific comments below are all I would ask for in terms of changes to the paper. However I would like to suggest an extension for your future work. The focus on biodiversity would I think be more than enough for most referees and readers. However, my personal view is that it is in principle just as susceptible to the criticism of single focus decision making as could be applied to disastrous policies such as the Common Agricultural Policy (certainly for the first three decades of its operation). The real world is not single focus. It is the meeting of two hugely complex systems, one environmental the other human. Each system is linked to the other in numerous, generally unplanned ways. This means that a single focus policy tends to have massive unforeseen consequences. The authors might argue that a focus on biodiversity would generate unforeseen positive benefits and certainly some positive externalities are to be expected. But there are also very likely to be negative consequences as well. The 'Food Security' wing of the agricultural lobby deliberately overplay their hand at every opportunity – but as with most long standing critiques there is a grain of truth there.
Now the present paper does a lot better than most. The inclusion of the cost variable helps to bring agricultural externalities into the prioritisation (even if it would be a crude approximation in a cost-benefit assessment; but this is a prioritisation telling us which is best, not telling us if any of these projects are cost-beneficial). However, there are many other impacts that are excluded. The one that might be considered further concerns greenhouse gas emissions and storage. It seems to me that this might be included in a further extension of this analysis without too much difficulty. One might argue that the change in net greenhouse gases was linearly related to the change (positive or negative) on the area of woodland conserved or lost. One could use a simple average quantity for this or use values that vary by biomes. Perhaps more attractively the authors could use a standard social cost of carbon value to convert this into an economic value and roll this into the cost function. This would still preserve the simplicity of a prioritisation analysis based on the cost-effectiveness of biodiversity conservation and avoiding the problems of getting into a cost-benefit frame.

The introduction of a linear net greenhouse gas removal value function which will have constant marginal values will contrast nicely with the declining marginal value curve assumed for biodiversity. This will illustrate a further very important issue regarding optimisation (one which I feel justifies the separate second paper); namely that changing the scope of impacts considered in an analysis can substantially change its outcomes. It is perfectly plausible that this extension might produce a considerably different outcome. Taking the two papers together could produce a very nice demonstration of optimisation principles: (i) The present paper: spatial prioritisation radically improves the efficiency of available resources (ii) A follow up paper: improving the incorporation of the complexities of the real world within analysis can result in significant changes to that prioritisation and further improve resource efficiency. I don't feel this is a necessary requirement for the present paper but I hope the authors might find it an interesting idea for their future research.

RESPONSE: We thank the referee for this great suggestion. We fully agree with the importance of integrated assessments that consider the impacts on multiple objectives simultaneously. In fact, in much of our other current work with a wider set of colleagues we are pursuing a similar line of research to what the referee suggests where we analyze multiple outcomes of interest related to land use and land management including biodiversity conservation, greenhouse gas emissions and storage (CO_2, N_2O, CH_4), net returns to agricultural crop production, net returns to grazing, net returns to forestry, and water quality (phosphorus, nitrogen, sediment). But that is the subject of a different set of papers.

Specific Comments

<table>
<tr><td>Lines</td><td>Text</td><td>Review Comments</td></tr>
<tr><td>3</td><td>Optimization</td><td>While it not needed in the summary paragraph, it is worth making clear to the reader that there is not just one approach to optimisation. Approaches such as greedy algorithms and combinatorial optimisation can give significantly different answers, especially when applied to a large set of options such as global forestry protection. The approach adopted needs clarification and justification.</td></tr>
<tr><td colspan="3">RESPONSE: We note in the abstract that we use “a dynamic optimization approach” rather than “the dynamic optimization approach.” We also talk about “The optimal conservation strategy for this formulation…” to emphasize that the results depend on the specific formulation of the problem that we have set up. In the results and discussion section, and in the Supplemental Information, we describe how different objectives, methods, or data, affect the outcome. We also discuss the performance of different heuristics. We feel that we have clearly defined and justified the approach that we have taken as our main approach throughout the paper.</td></tr>
<tr><td>5-7</td><td>The optimal conservation strategy accounting for species richness, costs of conservation, existing degree of protection, and threat of habitat loss</td><td>There is a need to clarify to the reader that there is no universal optima; rather the optimal solution changes with every reformulation of the objective – and there are multiple defensible objectives. If we argue that the use of land cost proxies the forgone benefits of market production such as agricultural output then perhaps the most important omission is in terms of carbon storage (mentioned further below). Including this would considerably change optimal outcomes.</td></tr>
<tr><td colspan="3">RESPONSE: See prior response.</td></tr>
</table>

<table>
<tr><td>4 vs 8 vs 10</td><td>global priorities for expanding protected areas
vs. priority ecoregions where further forest loss leads to large reductions in species conservation
vs. conservation efforts</td><td>You switch from talking about optimising for expanding protection to discussing further forest loss and then switch to conservation efforts. These are three different types of action within 6 lines presented as if they are the same thing. But these are not the same thing and you would undertake these different objectives in alternative locations. This is rather confusing.</td></tr>
<tr><td colspan="3">RESPONSE: We have rewritten the abstract to improve clarity. The conservation actions that can be taken are to either protect existing unprotected forest (which prevents further deforestation) or to invest in reforestation. The new abstract is as follows:

“Ongoing deforestation poses a major threat to biodiversity. With limited resources and imminent threats, deciding when as well as where to conserve is a fundamental question. We use a dynamic optimization approach to identify sequencing of where to conserve plant species in 458 forested ecoregions globally over the next 50 years. The optimization approach includes species richness in each forested ecoregion, complementarity of species across ecoregions, costs of conservation that rise with cumulative protection in an ecoregion, the existing degree of protection, the rate of deforestation, and the potential for reforestation in each ecoregion. The optimal conservation strategy for this formulation initially targets a small number of ecoregions where further deforestation leads to large reductions in species and where costs of conservation are low. In later years, conservation efforts spread to more ecoregions, investing in both expanded protection of primary forest and reforestation. The largest gains in species conservation come in Melanesia, South and Southeast Asia, the Anatolian Peninsula, northern South America, and Central America. The results illustrate the potentially large gains in conservation that can be made with carefully targeted investments.
”</td></tr>
<tr><td>12</td><td>A strategy using conservation return on-investment while ignoring threat</td><td>This seems like a fourth action – which is now really confusing as is the concept of “conservation return on-investment while ignoring threat” which is not something I intuitively understand. The term ‘threat’ is also not defined here.</td></tr>
<tr><td colspan="3">RESPONSE: We have rewritten the abstract to improve clarity (see prior response).</td></tr>
<tr><td>14</td><td>equal fractional protection targets</td><td>I don’t know what this means – and given I have worked in this field for a while I think this might be a common reaction to this term</td></tr>
<tr><td colspan="3">RESPONSE: We have rewritten the abstract to improve clarity. We no longer include discussion of heuristics in the abstract. The discussion of heuristics occurs in the results</td></tr>
</table>

<table>
<tr><td colspan="3">where we provide a short (but hopefully clear) description of each heuristic. Definitions of heuristics are also presented in the figure legend for Figure 2.</td></tr>
<tr><td>11-14</td><td>A strategy using conservation return on-investment while ignoring threat is close to optimal, while focusing solely on species rich areas, low-cost areas, highly threatened areas, or equal fractional protection targets achieves only 5-47% of potential conservation gains.</td><td>I get the general gist of this but overall I don’t really know what this is trying to say. Even terms such as “potential conservation gains” are not self-evident. Can the summary paragraph be rewritten so that, as is required, it is understood by the general Nature reader.</td></tr>
<tr><td colspan="3">RESPONSE: We no longer include discussion of heuristics in the abstract. The discussion of heuristics now occurs only in the results section where we have more space to clearly define the analysis and results.</td></tr>
<tr><td>14-17</td><td>Including the potential for reforestation further emphasizes the importance of priority ecoregions, while including land costs that rise with land scarcity works in the opposite direction.</td><td>Taken at face value this sentence does not make sense. Read literally it says that including the cost of land reduces the importance of priority ecoregions.</td></tr>
<tr><td colspan="3">RESPONSE: We have rewritten the abstract to improve clarity and removed detailed discussion of changes in the objective function or approach to the results section.</td></tr>
<tr><td>1-18</td><td>Summary paragraph</td><td>Overall the summary paragraph needs a thorough rewriting to make it generally accessible, clear and unambiguous.</td></tr>
<tr><td colspan="3">RESPONSE: We have rewritten the abstract to improve clarity (see responses above).</td></tr>
<tr><td>24</td><td>virtually all major habitats have seen declines</td><td>Technically this cannot be true. What has happened is that natural habitats have been progressively replaced by developed habitats.</td></tr>
<tr><td colspan="3">RESPONSE: We now say “virtually all major natural habitat types have seen declines.”</td></tr>
<tr><td>27-29</td><td>Since 1990, 178 million hectares of forest have been lost, representing 4.2% of forests that remained in 1990.</td><td>Minor: Rewrite to only mention 1990 once.</td></tr>
<tr><td colspan="3">RESPONSE: We have rewritten the sentence: “Since 1990, deforestation has reduced forest area by 178 million hectares (4.2% of forest area)[5].”</td></tr>
<tr><td>51-52</td><td>We used data on 458 forest ecoregions based on The Nature Conservancy's (TNC) global terrestrial ecoregions</td><td>Please clarify the extent of the analysis – are you considering all terrestrial area? Given the uncertainty in the summary paragraph its not clear to me if you are considering all forests, all areas that are</td></tr>
</table>

<table>
<tr><td></td><td></td><td>or could be/where forests or some other area.</td></tr>
<tr><td colspan="3">RESPONSE: We have added a sentence at the beginning of the Data and Approach section to explain what we use as the extent of forest: “For each forest ecoregion, we used the area of forest remaining in 2018 as the initial forested area and forest area in 2000 as the maximum extent for reforestation[17].”</td></tr>
<tr><td>56-57</td><td>and the number of plant species (TNC). We used plant species data because plants are an important component of overall biodiversity</td><td>I am uncertain how this measure is working in practice. Is it as follows:
Suppose we have three areas, A, B and C with species as follow (trees used for illustration):
Area A contains oak, ash and willow Area B contains oak, ash and willow Area C contains elm and birch
So using the scoring procedure described in the paper the biodiversity measures are as follows:
A= 3 B =3 C= 2
Therefore it seems to me that the optimisation process will select one of A or B to start off with. Lets suppose it picks A
Then it will select B because 3>2
Finally it selects C
However if resource constraints force us to only choose to protect 2 sites then only A and B are chosen. This satisfies the optimisation rule but leads to a poorer outcome because these two sites only protect 3 species rather than the 5 that would be protected is A and C were chosen.
Have I misunderstood what is happening here? If so then this needs clarifying. If not then the present approach needs justification.
Later on (line 85) you state that the approach is “implicitly assuming that all species are endemic”. I think this is perhaps how you avoid this problem – effectively by assuming that the commonality of species in Sites A and B (above) can never happen. Is this correct? If so then perhaps you should discuss the challenges that arise when that endemism assumption does not hold.</td></tr>
</table>

<table>
<tr><td></td><td></td><td>I'm not an ecologist so I can't judge how realistic the endemism assumption might be in a world split into just 458 forested Ecoregions – maybe that's perfectly defensible.</td></tr>
<tr><td colspan="3">RESPONSE: Following suggestions from other referees, we account for complementarity across ecoregions. Under complementarity analysis, in your example, if site A is already conserved, and oak, ash, and willow are conserved, then conserving site B will add no additional species. Conserving site B, given that site A is conserved, is 0. Conserving site C would add two additional species. So, if you could fully conserve two sites, you should choose either A and C, or B and C (but not A and B).

Also following suggestions from other referees, we have replaced what was our "base case" in prior drafts with the version of the model that incorporates complementarity from the beginning. We no longer assume endemism in the main text (which isn't realistic at the level of 458 forested ecoregions), though we examine the case of endemism in the SI to highlight the role of complementarity in our main results.</td></tr>
<tr><td>65-66</td><td>The number of species protected in each ecoregion is a power function of the forest area protected in that ecoregion ("species-protected-area curve").</td><td>This is defensible as far as it goes but I would like to see more self-critique. If you are optimising biodiversity and defining that in terms of the number of species conserved then you are ignoring abundance. Does this means that a state in which there are just a small number of each of X species will be preferred over a great abundance of X-1 species? If so then I think this needs clarifying as it would be debateable whether such an objective is what we want to optimise. Also, if we are optimising something that follows a power curve how does this relate to other benefits from the land which follow different relationships? For example, agricultural produce might be best approximated by a linear function. It seems plausible that this is captured in the linear cost function. But what about carbon storage? This might also be best approximated in a global analysis by a linear relationship with area – but its non-market value won't be captured by the GDP proxy. Does this matter? Perhaps not because the exercise is looking for a prioritisation i.e. it is answering a question about which order some pre-determined action should be taken in, rather than the question of whether action is actually</td></tr>
</table>

<table>
<tr><td></td><td></td><td>justified. However I would really like to see some discussion of these issues.</td></tr>
<tr><td colspan="3">RESPONSE: The power function relationship generates declining benefits as a function of remaining forest area. This along with (equal or) rising marginal costs in the amount of protected area means we have declining return on investment in conservation within an ecoregion. Including a benefit like carbon sequestration/storage, which is a function of what is done at a particular place (no complementarity or other spatial interaction effects), is straightforward. It will simply add to (or potentially subtract from) the net benefits of conserving in a particular place. As this isn't the focus of this paper, we didn't include this discussion in the paper. In addition, the species-area curve is probably the single most widely used and well understood function in conservation biology. We don't think our discussion of the strengths and weaknesses of this approach would add anything to what isn't already widely known.</td></tr>
<tr><td>73-74</td><td>Additional species protected per dollar spent</td><td>I refer back to my comment on lines 56-57; are these actually additional distinct species or more of the same but in different places? I really like this paper and don't want to tie it down to the point that its important message is lost but clarification of what is being done would help</td></tr>
<tr><td colspan="3">RESPONSE: We now incorporate complementarity from the beginning so that these are additional species not already conserved elsewhere.</td></tr>
<tr><td>88-89</td><td>We account for complementarity by developing a measure of species overlap among ecoregions.</td><td>Excellent – that addresses my concern above. I will leave those comments in just so you can see what readers reactions might be. You might think about heading those concerns off earlier just by flagging up that you test assumptions subsequently.</td></tr>
<tr><td colspan="3">RESPONSE: We think focusing on the case with complementarity from the beginning makes the presentation clearer and hopefully eliminates confusion associated with the previous version of the paper. We thank the reviewer for their comments that have helped us to improve the paper.</td></tr>
<tr><td>84-97</td><td></td><td>This paragraph did a lot to allay my concerns about the paper</td></tr>
<tr><td colspan="3">RESPONSE: We think presenting the case with complementarity, rising marginal costs, and reforestation potential from the beginning makes for a better presentation.</td></tr>
<tr><td>102-106</td><td>The largest gains in species conservation come in Melanesia, Southeast Asia, parts of China, the Anatolian Peninsula, northern South America, and Central America (Fig. 1). In</td><td>This makes a lot of sense as a result – but what is the relative strength of the cost variable in driving this result as opposed to perhaps higher biodiversity scores in tropical zones?</td></tr>
</table>

<table>
<tr><td></td><td>contrast, little forest is protected in many higher income regions with high land costs including the USA, Europe, Japan and Eastern China</td><td>A separate point is that many of these high income/high cost countries are pushing afforestation targets for greenhouse gas removal. A common critique of this is that it would be more cost-effective for those countries to pay for afforestation to happen in lower income countries s well. Does your result suggest that there might be a biodiversity bonus to such outsourcing of afforestation? (it seems plausible to me).</td></tr>
<tr><td colspan="3">RESPONSE: The results are driven by what ecoregions have a high conservation ROI. Since ROI is a ratio of marginal benefit to marginal cost, having a high ROI can occur with high marginal benefit and low marginal cost, or really high marginal benefit and high marginal cost, or low marginal benefit and really low marginal cost. Examples of each of these outcomes are shown in Figure 2a. For example, the East Deccan dry-evergreen forests (one of the circles in the upper right of Figure 2a above the dotted line) have very high marginal benefits and high marginal costs, while the Central Congolian lowland forests (one of the circles on the far left above the dotted line) have moderate marginal benefits but very low marginal costs.

Our results are fully consistent with the desirability of having high-income countries pay for afforestation (or better yet having them pay to prevent deforestation) in low-income tropical countries. There does tend to be high ROI in low-income tropical countries both because of high marginal benefits (species rich tropical forests) and low marginal costs.</td></tr>
<tr><td>110-111</td><td>In virtually all priority ecoregions (67 of 68), all land is either conserved or deforested within 10 years.</td><td>Does this imply that without action all of this land will be deforested in 10 years?</td></tr>
<tr><td colspan="3">RESPONSE: The combination of protecting remaining forest and on-going deforestation will mean that in some ecoregions there is no remaining unprotected forest. For example, in our analysis for the case with complementarity, rising costs, and reforestation we find that in 13 ecoregions all initial forest will either be protected or deforested by the end of year 1. This does not imply that in those same ecoregions, all land would be deforested in the absence of protection. In fact, there are only 11 regions in which conservation occurs where, in the absence of protection, all initially intact, unprotected forest would be deforested by the end of the planning horizon (these can be identified using spend and deforestation rates in Table S3 in the SI).</td></tr>
<tr><td>118-119</td><td>the rate of deforestation (threat)</td><td>This definition of threat could usefully have been clarified earlier.</td></tr>
<tr><td colspan="3">RESPONSE: The opening line of the abstract states that “Ongoing deforestation poses a major threat to biodiversity.” We now define threat in the third sentence of the Data and Approach section: “We defined threat for each ecoregion as the annual rate of deforestation (as a fraction of remaining unprotected forest), which is calculated as the annual rate forest cover loss between 2000 and 2018.”</td></tr>
</table>

<table>
<tr><td>129-
130</td><td>the benefit of spending the 50-year budget up front. The gains from immediate access to the entire budget are modest, an additional 124 species (0.12%).</td><td>This is a surprising result – what is driving this? It seems at odds with the earlier result at lines 110-111 which suggested to me that timing really mattered. How can it be that, without conservation, all priority ecoregions that are not already protected disappear – yet making a 50 year budget all available in year 1 makes so little difference? Some discussion of how this apparent contrast arises would be helpful.</td></tr>
<tr><td colspan="3">RESPONSE: We agree. We were surprised by this result as well. Our explanation of this result is as follows: “The gains from immediate access to the entire budget are modest, an additional 61 species (0.26%). Budget size and timing matter to the extent that they allow conservation to proceed before deforestation, but allocating a sufficiently large annual budget optimally captures much of this potential, leaving little remaining gain from shifting spending earlier in time.”

While we lack space to expand further in the text, our annual budget is large enough and many deforestation rates are low enough that, in ecoregions where we run out of intact forest to conserve, it is primarily due to successful conservation rather than deforestation. The small gains from having the budget up front reflect the modest loss of intact forest in those ecoregions under the optimal dynamic solution with an annual budget. We expand on these points in Section S1 and Figure S1 of the Supplementary Information. In the last paragraph of Section S1 we give the following explanation:

“Figure S1 also offers intuition as to why threat plays a more modest role in our empirical results and, as a result, why the cost-benefit heuristic performs so well. In our setup, deforestation (threat) serves only to constrain forest protection. The dynamic nature of our optimal solution will improve upon the cost-benefit heuristic when that constraint binds, i.e., when unprotected forest land runs out before the end of the planning horizon. Many ecoregions have a sufficiently large amount of unprotected forest area or high land costs that the constraint imposed by deforestation will not bind. In addition, even among ecoregions where the deforestation constraint binds, the dynamic solution will only differ from the cost-benefit heuristic in ecoregions in which the initial conservation ROI falls in a particular range. Specifically, threatened ecoregions in which the dynamic solution matters are those with an initial conservation ROI that is lower than in other areas, but where the ending ROI is at least as high as it will be in other areas. Those conditions define ecoregions where the cost-benefit heuristic would not go to initially (due to initial conservation ROI), but where not going there would result in a loss in conservation ROI later on (due to the combination of threat and ending ROI). This is precisely the type of area depicted by ecoregion C in Figure S1. Because this combination of features must align for threat to affect the optimal solution, we find only a modest role of threat and a modest benefit of a dynamically optimal solution (vs a cost-benefit heuristic) as applied to our dataset.”</td></tr>
<tr><td>140</td><td>the heuristic targets equal fractional protection</td><td>I did not understand this phrase</td></tr>
</table>

<table>
<tr><td colspan="3">RESPONSE: We apologize for being opaque on the prior draft. Some conservation proposals, like “30-by-30,” implicitly assume that the objective should be equal protection in all ecoregions (achieving 30% protected area by 2030). Since this is a common conservation approach we thought we should evaluate it. We find that it performs quite poorly. We would expect poor performance from an equal protection goal since it doesn’t take account of differences in biological benefits, economic costs, or any other factor that distinguishes ecoregions.</td></tr>
<tr><td>147-160</td><td></td><td>Very nice sensitivity analyses yielding expected but pleasing results.</td></tr>
<tr><td colspan="3">RESPONSE: Thank you.</td></tr>
<tr><td>212-222</td><td></td><td>The discussion regarding the incorporation of climate change impacts differs from the point I raise regarding net greenhouse gas removal (above).</td></tr>
<tr><td colspan="3">RESPONSE: We agree that this is a different point; please see our response to your earlier comment that referenced greenhouse gas removal.</td></tr>
<tr><td>235-238</td><td>we have not incorporated heterogeneity within ecoregions that could allow for targeted conservation in areas with greater biological benefit, lower cost, or more immediate threats. As better data becomes available, the analysis can be redone to refine our results.</td><td>Good to see this. Perhaps an earlier note regarding this would be helpful in clarifying the approach taken, its strengths and limitations.</td></tr>
<tr><td colspan="3">RESPONSE: We now include a short call out when we introduce the species-area curve relationship in the Data and Approach section: “The number of species protected in each ecoregion is a power function of the primary forest area protected and reforested area (“species-protected-area curve”), which assumes that land is homogeneous within an ecoregion.”</td></tr>
<tr><td>274-275</td><td>Figure 2. A) Relationship between the number of species protected in the first 10 years in each ecoregion (circle size)</td><td>Suggest this is rewritten as:
Figure 2. A) Relationship between the number of species protected in each ecoregion (circle size) in the first 10 years</td></tr>
<tr><td colspan="3">RESPONSE: We have rewritten the figure legend along the lines suggested: “Relationship between the expected number of species protected in each ecoregion (circle size) during the planning horizon…”</td></tr>
<tr><td>277-278</td><td>Threat is defined as the fraction of remaining unprotected forest that is deforested per year at current rates.</td><td>I would have appreciated seeing this definition earlier on in the paper.</td></tr>
<tr><td colspan="3">RESPONSE: Per our response to your comment above on what were lines 118-119, we now include this definition of threat as the third sentence of the Data and Approach section.</td></tr>
</table>

<table>
<tr><td>296</td><td>Figure 3</td><td>The small size of these four global maps makes the information difficult to read. First impressions are that they are all very similar. Looking at them in the unconverted form the differences are much clearer – but that is in major part because the figures are simply much bigger there. This is a production issue (rather than one for the authors) but excessive reduction will render these figures unreadable.</td></tr>
<tr><td colspan="3">RESPONSE: We moved Figure 3 to the Supplementary Information. We did this mostly because these results were less important with the rearrangement of the presentation to include complementarity, rising costs, and reforestation from the beginning. We also agree with you that presenting four panels makes the results too small to really see. Moving this to the Supplement allows us to not worry about space constraints and present each map separately so that it is big enough to be seen.</td></tr>
<tr><td>43-45 vs 452-475</td><td></td><td>On 43-45 the paper states “Here we apply a dynamic optimization approach to maximize the number of plant species conserved in the world’s forests” but 452-475 discusses “the abundance of profiles for mammals across the same ecoregions”. I found this confusing as I had assumed throughout my reading of the paper up to that point that the analysis was solely concerns with plants and that the concept of complementarity was also in terms of plants. I admit I am confused by this.</td></tr>
<tr><td colspan="3">RESPONSE: We have thoroughly rewritten the Methods section. Incorporating complementarity for our data, which does not have individual species listed in each ecoregion but only the overall number, required a different approach, which we now lay out in detail. We think the rewritten methods section provides a much cleaner and clearer exposition.

In short, we lack what we consider to be adequate plant species lists by ecoregion to assess complementarity species-by-species; we have only species richness data. To overcome this limitation we use the degree of overlap in mammal species lists across ecoregions as a proxy for the degree of overlap in plant species lists across ecoregions. Our probabilistic framework relies on the degree of overlap only and can be reapplied as improved plant species lists become available.</td></tr>
</table>

Referee #4 (Remarks to the Author):

Your efforts in sorting out priorities for conservation of tropical forests are original and well worth pursuing. The quality of data a method of analysis are complete and compelling. The conclusions seem to me to be robust and very useful in this point of overall analysis of evidence in this situation. I have no further references to suggest and no negative comments on clarity.

RESPONSE: We thank the reviewer for their comments that have helped us to improve the paper.

Referee #5 (Remarks to the Author):

I have re-read the manuscript and all responses to reviewer comments. The authors did a very thorough job of responding to reviewer comments, and this version of the paper is a more robust contribution. I especially appreciate the authors considering my comments to add complementarity, including new analyses and results. I believe this will be a strong contribution to the global conservation literature, and I have no further comments.
Laura Pollock

RESPONSE: We thank the reviewer for their comments (especially pushing us to include complementarity) that have helped us to improve the paper.

Reviewer Reports on the Second Revision:

Referees' comments:

Referee #1 (Remarks to the Author):

I appreciate the authors responses to my comments and I have not further comments on the manuscript.

Referee #2 (Remarks to the Author):

The authors have substantially improved the manuscript over the last two revisions. I have no further major concerns to raise about this work. My only remaining minor concern is that in addressing reviewer comments the Discussion often focuses on technical issues (e.g. model sensitivity issues), possibly at the expense of greater exploration of the policy implications of this analysis and the geographic insights that emerge from the analysis. For example, Fig 1 clearly indicates that most priorities are in the less wealthy tropical countries and few in wealthy temperate countries, implying that innovations in global conservation funding mechanisms are urgently required. Conservation finance is a major barrier to implementing this sort of plan. How do we begin to solve that? Are international conservation finance mechanisms high enough on the agenda with the CBD and Post-2020 GBF? Even more problematic is that some existing major international finance mechanisms relate to development (e.g. see Bill Laurence's papers on how China's development funding is detrimental to conservation in several developing countries). So not only is there is a woeful shortfall in effective international conservation finance, it is also directly counteracted by international development finance in some cases.

Another example of where greater focus on the geographic aspects of the results could be developed is to highlight the developed countries that appear to be failing with respect to conservation. Fig 1 seems to indicate that Australia and New Zealand are the only two developed countries with yellow circles indicating "ecoregions in which all initial forest is conserved or deforested in the first year". Does that not indicate a gross inadequacy of environmental management in these systems that should be called out?

The priority ecoregions in Fig 1 often appear to occur in countries with low governance scores. There is a fundamental problem with this: issues like poor governance have created the very circumstances that result in these geographies being selected as conservation priorities, yet those issues are likely to also be significant barriers to achieving good conservation outcomes in the future. Is the implication of this analysis that we either have to solve that fundamental problem or risk suffering major conservation opportunity costs?

The authors may wish to consider shifting a portion of the discussion away from more technical considerations to further explore the broader issues highlighted above in order to further improve the potential for impact of the paper, though I do not consider this a requirement.

Congratulations on authoring an interesting and insightful paper.

Referee #3 (Remarks to the Author):

Review of further revised MS:
"Where to protect first in global forest conservation"
by Ian Luby, Steve J. Miller, Stephen Polasky
Manuscript number 2021-02-01982A-Z

Thank you for the opportunity to re-review this revised paper.

General Comments

The paper highlights the importance of trying to focus conservation resources upon those priority areas where they will be most efficiently applied, generating the greatest impact. The latest revision has fully addressed the major problems of the previous version of this manuscript. I have one very small suggestion for further change and whether this is responded to or not, I support publication.

Specific comments

The authors have responded positively to all my comments. I particularly appreciate their decision to drop the base line model and focus on the model including complementarity. This feels much more defensible from the outset. I'm sorry that this clearly necessitated a considerable degree of rewriting though I can see that another Reviewer effectively asked for the same change.

I also greatly appreciate that the authors attended to the large number of text changes I requested including the rewriting of the Abstract. I feel this removes the unevenness of the writing style in the previous manuscript.

I am glad to hear that the authors feel the suggested developments to this paper are useless and that they have already been working on such extensions.

I have one small suggestion for a further change which the authors might consider:

Replace:
222-3:
spatial heterogeneity within ecoregions, potential tipping points, and ….

With:
spatial heterogeneity within ecoregions, potential tipping points across ecoregions, and ….

The following recent paper is highly relevant here:
Boulton, C.A., Lenton, T.M. & Boers, N. Pronounced loss of Amazon rainforest resilience since the early 2000s. Nature Climate Change, 12, 271–278 (2022). https://doi.org/10.1038/s41558-022-

01287-8

I suggest either adding this to or replacing the existing citations in this sentence:

While this is only a small change, it more explicitly acknowledges the problem of resilience loss and consequent tipping point effects across ecoregions, which in turn would alter the prioritisation ordering for conservation.

Irrespective of whether the authors accept this suggestion, I am happy to recommend the present manuscript for publication and congratulate the authors on a high quality contribution to the literature.

Ian Bateman