COPENHAGEN CONSENSUS CENTER

A Perspective Paper on Climate Engineering as a Response to Climate Change

Anne E Smith

COPENHAGEN CONSENSUS CENTER

COPENHAGEN CONSENSUS ON CLIMATE

A Perspective Paper on Climate Engineering as a Response to Climate Change

Anne E Smith

Charles River Associates
1201 F Street, NW, Suite 700
Washington D.C. 20004

Copenhagen Consensus Center
Copenhagen Business School
Solbjerg Plads 3
DK-2000 Frederiksberg
Denmark

+45 3815 2255
info.ccc@cbs.dk
www.copenhagenconsensus.com

PREFACE

ABSTRACT

In their Assessment Paper, "The Potential Benefits and Costs of Climate Engineering: A Case for Research," Bickel and Lane have provided a deterministic case for funding a long-term, intensive research and development (R&D) program for geoengineering. Because their estimate of the requisite R&D budget is about 0.1% of their deterministic estimates of the net benefits of using geoengineering, they argue that it is unlikely any of the uncertainties they did not analyze would affect the case for performing the R&D. This Perspectives Paper overlays a consideration of potential unintended consequences from geoengineering onto their analysis and extends it with calculations of value of information from the R&D. It finds that the value of ***perfect*** information is indeed much higher than Bickel and Lane's proposed research budget for almost all but the most extreme assumptions about probabilities (either optimistic or pessimistic) on whether geoengineering will produce significant unintended consequences. However, it also finds that ***imperfect*** information may have zero value for a wide range of assumptions. Thus, a standard analysis of value of information seems to undercut a view that uncertainties not addressed in the Assessment Paper are unlikely to affect the merits of conducting the further research on geoengineering. This Perspectives Paper, however, also takes a more critical look at the theoretical assumptions underpinning the standard formula for value of information, and finds that they may be inappropriate in a public policy making process. This paper suggests an alternative value of information formula to match the current issue's role as part of societal decision making by groups who hold very different sets of probability assumptions. When the proposed alternative calculation of value of information in a social choice context is applied, one obtains much larger value of information estimates over a broad range of probability assumptions.

COPENHAGEN CONSENSUS ON CLIMATE

The Copenhagen Consensus Center has commissioned 21 papers to examine the costs and benefits of different solutions to global warming. The project's goal is to answer the question:

"If the global community wants to spend up to, say $250 billion per year over the next 10 years to diminish the adverse effects of climate changes, and to do most good for the world, which solutions would yield the greatest net benefits?"

The series of papers is divided into Assessment Papers and Perspective Papers. Each Assessment Paper outlines the costs and benefits of one way to respond to global warming. Each Perspective Paper reviews the assumptions and analyses made within an Assessment Paper.

It is hoped that, as a body of work, this research will provide a foundation for an informed debate about the best way to respond to this threat.

CONTENTS

INTRODUCTION

The Assessment Paper of Bickel and Lane provides a thorough summary of the range of options that fall under the term "geoengineering," and some detail on the specific technological concepts for options that fall under the classification of solar radiation management (SRM). It provides less information on technological proposals for accomplishing air capture (AC), but does at least explain how SRM and AC would differ in the way they might help address risks of climate change. The paper is clear that it focuses on providing a first-cut assessment of the potential costs and benefits of these general categories of climate engineering, relying on a limited body of existing literature. It does, nevertheless, develop its own new estimates of the benefits of SRM and AC using the DICE model. Overall, this paper is an excellent review and analysis for anyone who wishes to gain a high-level understanding of how geoengineering methods could become an important element of long-term management of global climate change. However, what becomes apparent throughout the paper is the large range of risks and unknowns associated with any of these methods, but most particularly with SRM. Even though the Assessment Paper's scope was necessarily limited to providing a few point estimates of benefits and costs, and the benefit-cost ratios, the issue of uncertainties could have been addressed more satisfactorily without resorting to a full-blown "probabilistic analysis." In this Perspectives Paper, I attempt to show how one can extend the initial efforts of Bickel and Lane to more directly assess the implications of these enormous uncertainties on their conclusions.

Bickel and Lane have performed a deterministic benefit cost analysis for SRM and AC. A range of cost estimates is noted in the discussion, some of which are treated with an abbreviated form of sensitivity analysis. Uncertainty in the benefits is, of course, also enormous, and not explored through sensitivity analysis. Rather, the deterministic analysis produces such huge benefit-cost ratios, and huge net benefits as well, that the authors conclude that the very large set of unknowns and uncertainties in these estimates merit investment in a concerted program of research and development (R&D). Their proposal for that R&D budget is then developed from a bottom-up process of thinking through the steps and elements of such a program. Their recommendation is $0.75 billion dollars per year,[1] although they describe an R&D program that would start much smaller and expand rapidly as the time of full deployment would approach.[2] It is useful to restate their proposed R&D budget as a present value so that it can be more readily compared to their net benefits estimates, and to the estimates of value of information that I will be reporting below. I estimate their recommended R&D spending to have a present value of about $5 to 10 billion.[3] This contrasts to their deterministic net benefits estimates that range from $4 trillion to $25 trillion.[4]

1 Bickel and Lane.

2 "As R&D progresses, and assuming that results were favorable, spending would increase from tens of millions of dollars in early years to the low billions of dollars." (Bickel and Lane, p. 6.)

3 I based this on an assumption that the spending that they describe will extend through 2025, given that they have suggested that 2025 would be the earliest reasonable date when the geoengineering technologies would be ready for full deployment if R&D were to start now (Bickel and Lane, p. 20). I applied a discount rate of 5.5%, as the authors have been using to several spending patterns from 2010 to 2025 that have an average spending level of about $750 million per year.

4 These are the range of benefits estimates for SRM 1 through SRM 3, under a variety of start dates, discount rates, and mitigation policies, presented in Table 7 of Bickel and Lane (p. 42) minus the range of respective costs in their Tables 9 and 10. It includes net benefits of AC, but they estimate AC's net benefits at about zero (i.e., have a maximum net benefit ratio of 1).

Bickel and Lane do not advocate actually initiating SRM or AC until the R&D has proven that the uncertainties not addressed in their deterministic analysis will not eliminate the apparently very large net-benefits case for climate engineering. They even imply that a mature SRM and AC capability might never be deployed, but only that the concepts need to be developed into real options that the world can resort to if climate changes turn out to be far worse than in their own benefits analysis.

While these arguments appear to make a good case for a spending on R&D, the failure of the Assessment Paper to address the role of the uncertainties in this decision to commit to creating a real option makes the case incomplete. In fact, the existing professional and public pressure against even a modest research program around climate engineering is itself tied to the exceptional nature of the associated uncertainties. There are uncertainties about whether it will work, what it will cost to do, and what benefits it might produce in terms of climate damages avoided. But the most important uncertainty at play in this debate is concern about unleashing an entirely new set of environmental or economic damages in the course of trying to reduce the costs of global warming. Discussion of these risks tends to be non-specific, and lumped into a single ill-defined category called "unintended consequences." Even though they are ill-defined, these risks are clearly central to the debate about whether to fund even R&D for geoengineering, and they need to be addressed head-on in an analysis that is intended to make the case for such R&D.

INCORPORATING CONSIDERATION OF UNCERTAINTIES WITH VALUE OF INFORMATION ANALYSIS

Most R&D investment decisions are made in the face of a similar array of uncertainties, and a standard method for assessing whether and how much to spend to obtain more information is "value of information" analysis. Value of information analysis directly estimates how having better information before taking action can improve the quality of future decisions – in this case, whether to deploy climate engineering options in 2025 or later. If research has a chance of identifying an undesirable cost of climate engineering before full scale deployment occurs, it can help us avoid using climate engineering if it would reduce rather than increase net benefits. Research can therefore increase the expected benefits of climate engineering by cutting off the downside risks that are presently associated with this option for climate risk management, while maintaining its upside potential. It does so by directly considering the extent to which the future deployment decision could be altered by resolving any of the many uncertainties before making that decision. The value of research is determined by comparing the expected net benefits of the climate engineering decision when taken in the face of present uncertainty to the expected net benefit of the same decision when taken in the face of less uncertainty. If the latter is larger, the information has value. The difference between the two expected net benefits is viewed as the most that one should be willing to pay in order to perform the research that would provide that reduction in uncertainty. If the R&D program to obtain that information would cost less than the assessed value of information, then it is a good investment. If it costs more, then the best course of action is to make the decision under uncertainty, without the additional information-gathering. In the context of the geoengineering assessment, if the value of information for the R&D program outlined by Bickel and Lane exceeds $10 billion, one has a case for accepting their proposal.

Bickel and Lane's initial deterministic estimates of the net benefits of SRM range from about $4 trillion to $25 trillion (present values), with benefit-cost ratios all exceeding 20. They argue that despite the significant uncertainties they did not analyze, it is difficult to imagine any situation in which an investment of $10 billion (present value) would be a bad allocation of resources. However, value of information analysis can show otherwise - all depending on initial views (assumptions) about the possibility of significant unintended consequences, and their potential magnitude. For this Perspectives Paper, I have developed a fairly simple structure for estimating value of information for this particular issue, and I have used it to prepare estimates of value of information based on the net benefits calculations in Bickel and Lane. Because value of information is contingent on current views about the potential size and probabilities of unintended consequences, I do not provide a single estimate of "the" value of information, but instead show how that value varies over the full range of possible viewpoints. These current viewpoints are known in technical terms as the "prior probabilities" regarding the uncertainty that is to be studied. The prior probability is a subjective assumption, and thus it can vary widely among individuals. Thus, some people may see great value in an R&D program that others would see as having no value at all, depending on their respective "priors."

In the following, I simplify the problem to consider only one crucial uncertainty, which is the possibility that SRM can lead to serious "unintended consequences" that would add new environmental costs of their own, and exacerbate climate damages as well due to sudden more rapid climate change when the program of SRM is suddenly stopped. Consideration of this single uncertainty creates some interesting patterns in the potential value of the proposed R&D program that may provide a platform for a much richer discussion of its merits. The proposed R&D program would presumably reduce uncertainties over a variety of other uncertainties such as cost, functionality, climate change damages, etc., and thus value of information from this single program might be larger or smaller than the estimates I will provide here. While I doubt they would significantly alter those patterns, a more comprehensive analysis may be desirable. I hope that the initial structure that I present in this paper can serve as a foundation on which that more comprehensive discussion and analysis can be built.

I also wish to note that my example, like the Assessment Paper, focuses on SRM. However, the Assessment Paper does recommend research on AC in addition to SRM. This recommendation may seem a bit weak because Bickel and Lane conclude that AC may have a benefit-cost ratio less than one, which is not only low, but much less than they estimate for SRM.[5] However, I would like to emphasize that the AC option not only has gross benefits as high as $5.5 trillion,[6] but it has risk management properties that could easily lend it a large positive value of information even if it does have negative net benefits when calculated deterministically as Bickel and Lane have done. Indeed, there may be some interactions between the value of information for AC and SRM when considered as a set of options rather than as individual options. While I have not extended my analysis to AC in this Perspectives Paper, I consider that an important next step.

5 Bickel and Lane report that the benefit-cost ratio from their deterministic estimates for AC are only 1 or less.

6 Bickel and Lane, Table 6, p. 41.

STRUCTURING THE SRM ISSUE AS A DECISION MADE WITH UNCERTAINTY

Figure 1 illustrates the essence of the deterministic analysis in Bickel and Lane using a decision tree diagram. The numerical values on which I will base my value of information analysis are for the SRM 3 option and for the "temperature constraints" policy case.[7] The costs of each option are shown as the outcomes of each decision branch. Climate damages and mitigation costs of the option come from Table 4 of the Assessment Paper, while costs of the SRM (on the SRM decision branch only) are from Table 10. The total cost of each decision is the sum of these three categories, and is $4.9 trillion for the SRM option and $22.8 trillion for the No SRM option (i.e., "SRM 0" in the Assessment Paper). Clearly the least-cost decision is SRM, and the savings of having SRM as an option in this decision tree is $17.9 trillion (i.e., 22.8-4.9). This result is fully consistent with the Assessment Paper's benefit-cost ratio of 27 for SRM 3, but this analysis is computing values as differences in total costs of two options, rather than as differences in climate damages plus mitigation costs (i.e., (11.9+10.9)-(3.6+.6) = 18.6) divided by differences in SRM costs (i.e., 0.7).

Figure 2 carries the analysis one step closer to the actual situation that characterizes the decision of whether to do SRM or not. If No SRM is undertaken, the same outcomes occur as for that branch in Figure 1. However, the decision branch for SRM now reflects the concern with "unintended consequences" that Bickel and Lane describe but do not account for in their deterministic benefit-cost analysis. Figure 2 presents the SRM decision as one that is taken without any advance investment in research on this risk. That is, Figure 2 accounts for the key uncertainty in the SRM decision, but does so as if nothing can be learned from R&D prior to making that decision. Instead, if one proceeds down the SRM decision path in 2025, there is some probability, p, that SRM will not cause any unintended consequences at all, and its costs will be just like Bickel and Lane estimated (i.e., the same as for that branch in Figure 1).

However, Figure 2 shows that there also is a probability, (1-p), that unintended consequences will start to become apparent. At that time (e.g., 2035 in this example) the SRM will have to be stopped abruptly, costs of climate mitigation will increase to the level in the No SRM case (or perhaps even higher because of the need to act even more rapidly on those controls than if they had been initiated much earlier), climate damages may rise substantially, due to a more rapid rate of change when accumulated CO_2 suddenly is allowed to have its full warming potential, which had been offset by the SRM for many years, *and* there would be the new costs from the bad consequences of the SRM itself. In return, the cost of the SRM would not be as high, as only its fixed costs and first ten years of operating costs would be incurred. In Figure 2, one can see my approximate estimates of how these various cost elements would change relative to the costs that were formally estimated by Bickel and Lane for the other two branches of the decision tree. These are rough but reasonable estimates for purposes of

7 I use their SRM 3 case as the only SRM option in the decision tree because it always has the largest net benefits, and any decision tree would thus always ignore the other SRM options even if they were included. I use the higher cost estimates of SRM associated with stratospheric deployment for these illustrative purposes. A complete analysis would include both forms of SRM as well as AC. I also focus on the case in which every climate management strategy must limit temperature change to 2°C, because it is easiest to estimates of alternative mitigation costs and climate damages under different outcomes than those actually presented in the Assessment Paper without actually re-running the DICE model, which was beyond the scope of what I could do for this Perspectives Paper.

the illustrative analysis. However, all but the SRM-related damages can be estimated precisely simply by running this additional scenario using the same version of DICE that Bickel and Lane have used.

When this extra branch is added, the decision of whether to undertake SRM is altered, and so too is the value of SRM. The calculations in Figure 2 can be run for any number of combinations of probabilities that SRM will not produce unintended consequences (i.e., will "be safe"), and any degree of damage from any unintended consequences. Table 1 shows a set of such results, including the case that reduces Figure 2 back to Figure 1, which is where the prior probability that SRM will be safe, p, is 1.0. Thus the deterministic result of Bickel and Lane has the maximal SRM value of $17.9 trillion that appears on the bottom row of Table 1, for p=1. However, as the prior probability p drops towards zero, the value of SRM drops. It drops more rapidly when the associated damage from the unintended consequences increases (i.e., as one moves to columns to the right in Table 1). The value of SRM is negative for

Figure 1. Representation of Deterministic Analysis of Assessment Paper as a Decision Tree

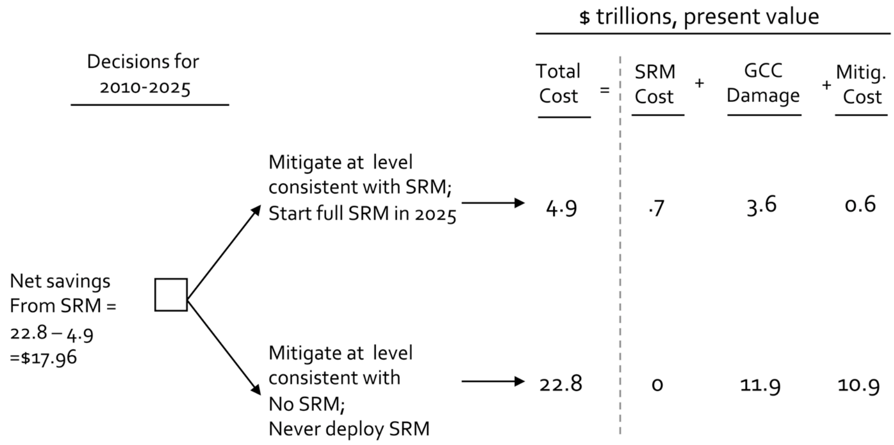

Figure 2. Representation of SRM Decision as a Decision under Uncertainty

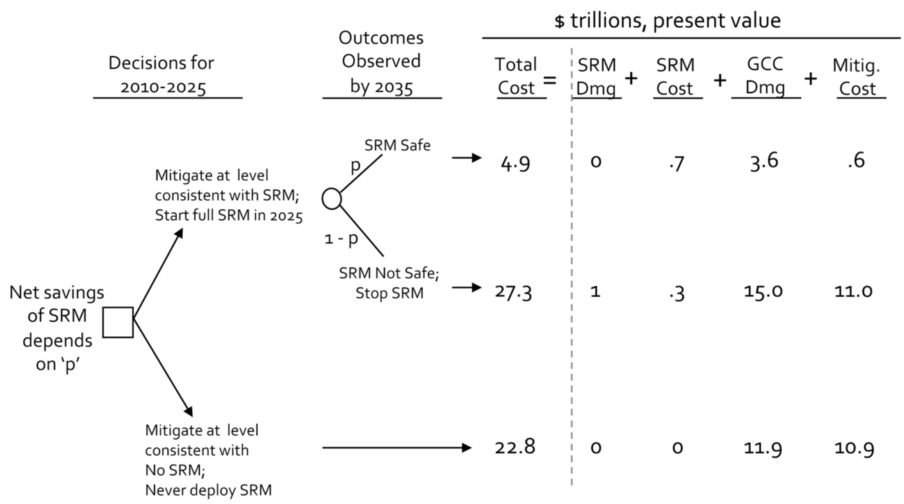

Table 1. Expected Cost Savings of SRM Strategy Compared to a No-SRM Strategy (Assuming no further research on SRM risks prior to SRM deployment decision)

		Damages from SRM if SRM not Safe (PV, $trillions)			
		1	5	10	15
prior probability	0	-4.5	-8.5	-13.5	-18.5
that SRM is safe, 'p' (subjective, varies with the person)	0.25	1.1	-1.9	-5.6	-9.4
	0.5	6.7	4.7	2.2	-0.3
	0.75	12.3	11.3	10.1	8.8
	1	17.9	17.9	17.9	17.9

those who have “priors”, or subjective views of the probability and the potential impact that make the ***expected*** cost of SRM higher than the $22.8 trillion cost of the No SRM option. The prevalent opposition to geoengineering that exists today may be traced to people whose individual “priors” place them in the part of Table 1 that has negative values. In those cases, the preferred decision in the decision tree is the No SRM option and, formally, the value of the SRM option is zero. However, the large negative values in Table 1 provide an indication of the degree of concern with expected damages that may be held by some people that SRM will be deployed anyway. They may have these concerns because they think that people who will be in control of the deployment decision may have the more optimistic views that lead to a large positive subjective value for SRM.

For those who believe that there is a better than 25% probability that SRM will not have unintended consequences greater than $1 trillion (present value), the decision to act on SRM has expected benefits exceeding $1.1 trillion. For those who believe that the direct damages of SRM could be about $10 trillion (i.e., could be larger than the damages projected for climate change itself), SRM would still have a large positive value even if one assigns that outcome a 50% probability. Nevertheless, one can also see that people who assign very high probabilities to the possibility of creating “another” problem of equal proportion to climate change would certainly disagree with a planned course of action that could lead to use of SRM.

Table 1 only presents the potential expected benefits of SRM, given this single large uncertainty. It may help clarify why there are such strong views both for and against SRM, given the current state of uncertainty. However, it does not reveal how a research program to better inform these decisions can reduce the risks of making a bad decision on SRM, while still preserving an option to use it if it proves to be as benign as assumed in the Bickel and Lane analysis. The value of information is the difference between the expected costs in the case where a research program is conducted before deciding whether to deploy SRM, and the expected costs in the case depicted in Figure 2 where the best decision is made given current information.

STRUCTURING THE SRM ISSUE TO ACCOUNT FOR INFORMATION FROM AN R&D PROGRAM IN ADVANCE OF ANY SRM DEPLOYMENT DECISIONS

Figure 3 shows how the decision tree is altered to estimate the expected costs when a research program is conducted before making any deployment decision. This figure has the same basic elements of Figure 2, except that the second stage of the decision process is preceded by new (but probably imperfect) information that would be produced by the R&D program. The decision today is now simply to choose the level of mitigation spending to engage in while the R&D on SRM is being conducted. Consistent with the previous examples, this current decision is simplified to include only whether to start with a level of mitigation that would be taken if one does ***not*** anticipate later having SRM as part of the management path, or with a mitigation level that would be consistent with SRM being deployed in 2025. Then, ***before the decision on whether to deploy SRM is taken***, new information is assumed to result for the R&D program. This information has been simplified to be either that the research finds that SRM will be safe or not safe. (In the figure, results of the research are identified by placing the outcome inside quotation marks. This indicates that they are merely conclusions, and not necessarily the actual state of the world.)

The decision on whether to deploy the SRM is taken after the research findings are reported, and thus is a more informed, less risky decision than in the case of Figure 2; that is, that decision is made ***contingent*** on the R&D findings. Thus, whatever one's ***a priori*** assumption about the probability of unintended consequences occurring from SRM (which decision analysts refer to as one's "prior"), if the research concludes "SRM will be safe," the probability that SRM is indeed safe will be increased. If the R&D has no risk whatsoever of coming to an erroneous conclusion, then the probability that SRM is safe becomes 1.0, regardless of any prior views, and the SRM deployment decision can be made in a risk-free manner. (This is the case used to compute the "value of perfect information.") However, most research efforts leave some chance of producing a false positive or false negative result, and the best that such research findings can hope to do is reduce but not eliminate the possibility of making an incorrect decision.

The probabilities associated with the branches of Figure 3 are derived using Bayes' Rule, and depend on the subjective prior probability, p, that was used in Figure 2 that unintended consequences will result from SRM. This probability is now also complemented by the probability that the R&D will produce either a false negative, q, or a false positive, s. The research produces a false negative if it concludes that "SRM is safe" even though SRM will not be found to be safe if actually deployed. Research produces a false positive when it concludes that "SRM is not safe" but it would actually be consequence-free if it were to be deployed. Each erroneous conclusion leads to a different potentially bad outcome. False negatives may lead to too much willingness to engage in SRM and false positives can lead to overly precautionary approaches that prevent society from benefiting from the cost reductions that SRM could otherwise provide.

As in the case of no R&D (i.e., Figure 2), the decision one takes, and the willingness to undertake SRM, is contingent on one's prior probability of the unintended consequences, but also on one's view of the ability of the R&D to properly identify the truth about those risks.

Figure 3. Decision Tree for SRM Decision Made with Improved Information from R&D

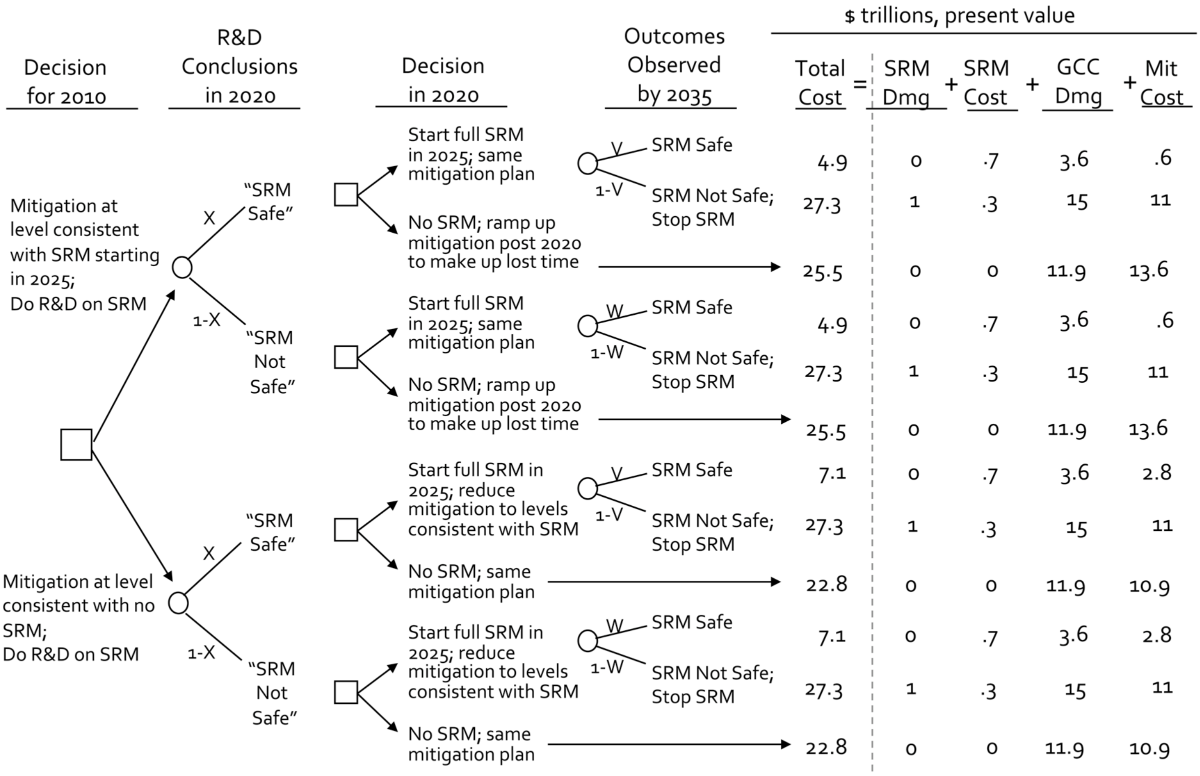

Notes to Figure 3:

X: Probability that R&D will conclude that SRM will not produce any damages of its own, i.e., R&D out come will be "SRM Safe." X includes the possibility that such a conclusion is a false negative.

$X = p(1-s) + (1-p)q$,

where:

p = prior probability that SRM will not produce damages of its own, i.e., pr(SRM Safe), as used in Figure 2; s is the probability the R&D will report a false positive, i.e., pr("SRM Not Safe"|SRM Safe) and q is the probability the R&D will report a false negative, i.e., pr("SRM Safe"|SRM Not Safe).

V: Probability that SRM is safe if R&D has concluded it is safe, i.e., pr(SRM Safe | "SRM Safe"). This is a posterior value of p, conditioned on one possible R&D outcome.

$$V = \frac{p(1-s)}{X}$$

W: Probability that SRM is safe even though R&D has concluded it is not safe, i.e., pr(SRM Safe | "SRM Not Safe"). This is another posterior value of p, conditioned on a different R&D outcome than in the case of V.

$$W = \frac{ps}{1-X}$$

Note on perfect information: R&D would produce perfect information as it has no probability of returning false positives or false negatives. This is the same as setting q=s=0. If this is done, X = p, and V=W=0. This reduces the above probability tree to a simpler tree used in calculating the value of perfect information, and this value of perfect information is also being presented in the results of this Perspectives Paper, for cases where both q and s are zero.

The latter are embodied in the probabilities of false negatives and positives. The value of information is the difference between the expected cost of the climate management problem as structured in Figure 3 and the expected value if that same risky decision is made without the benefit of first learning the research conclusions, which is the expected value from the decision tree in Figure 2 when using all the same parameter assumptions for p, and for the various outcome costs.

In order to perform any numerical computations using the value of information structure presented here, I needed estimates of a number of costs that I cannot obtain from the Assessment Paper, because they would require three additional DICE model runs. Because execution of those runs was not within the scope of this effort, I have made some educated guesses for those values in Figure 3. The basis for each of the cost estimate values in Figure 3 is explained in an addendum at the end of this paper. The salient points to note about them here are, first, that the values of SRM cost, mitigation cost, and global climate change damages were tied as closely as possible to the estimates of the Assessment Paper. The analysis can be fine-tuned later, if desired, while still relying on the basic analytic framework that I have prepared so far. Second, the present value for potential damages from SRM if SRM proves unsafe is set at $1 trillion, consistent with the example in Figure 2, and with the results in column 1 of Table 1. Value of information results presented below will change substantially if much higher or lower values are used for the likely magnitude of potential SRM damages, and sensitivity analysis on that assumption may be warranted and would be easy to do. However, for brevity, none are not presented in this paper.

VALUE OF INFORMATION RESULTS

Table 2 shows the value of information given a range of prior probabilities, p, that SRM will be safe, and a range of probabilities that the research will produce a false negative. All of the values in Table 2 assume that the probability of a false positive from the research is zero, and the costs if SRM is deployed but found to have unintended consequences are as shown in Figure 2 (i.e., a present value of additional damages being $1 trillion, plus additional climate-related damages and mitigation costs to catch up).

Table 2 shows that the value of information is not always greater than zero, despite the very large net benefits and benefit-cost ratio in Bickel and Lane. However, the conditions that lead to zero value of information only occur if one has a very pessimistic view (i.e., less than 25% probability) that SRM will not produce unintended consequences. Further, the value of information is well above $10 billion – the proposed R&D funding level – when it is positive. In fact, except for priors that SRM is more than 90% likely to be free of unintended consequences, the value of information is about $100 billion or more. At its peak, its value is in the range of $1 trillion to $3 trillion. (The value of perfect information is shown in the values on the leftmost column of Table 2, where the probability of a false negative is zero. Since this entire table assumes the probability of a false positive is zero, the information from research has its peak value, the "value of perfect information" in that column. That peak value is about $3 trillion.)

It is also interesting to note that the value of information is at its highest levels when one has only modest priors that SRM will be safe – where p is in the range of 10% to 50%. The most optimistic views are associated with declining value of information. One can also observe from Table 2 that value of information declines as confidence in the research to identify real problems declines. However, even if the research has a 75% probability of failing to detect a real problem with SRM, it still has value in the range of $100 billion to $200 billion (present value) if one's priors that SRM will be safe are anywhere in the range from 20% to 80%.

Table 3 shows the results for the same circumstances as in Table 2 except that the probability of a false positive has been set at 25%. In this case, the value of information is very widely zero. The remaining zone of positive value of information is greatly condensed around the range where it is at its peak in Table 2, and when in that range, the value still rises to about $2 trillion. However, if the research is likely to be relatively prone to concluding there will be problems with SRM-related damages when there will not actually be any, the people who would be most likely to find no value from R&D would be the optimistic individuals. The reason is that their priors that SRM will be safe are so strong that they will tend to interpret research results of potentially bad outcomes as more likely false outcomes than correct outcomes. They will be inclined to undertake SRM regardless of research results, and hence that research has relatively little chance of changing their course of actions.

When situations like that occur, research has no value. For people on the pessimistic end of the spectrum, the R&D has no value because it will not be good enough to convince people that geoengineering is safe, even if the research does not identify any concerns. For people on the optimistic end of the spectrum, any negative findings of the R&D would only confirm what they already believe, and they may tend to interpret any positive findings (red flags) as more likely being false positives than as information that might alter the preferred course of action with respect to using SRM. Research has to have the potential to persuade people towards different courses of action depending on its findings if it is to have any actual value. So, almost counterintuitively, we find that the very people who might be most inclined to undertake SRM are also the most likely to find no value in the research at all, if the fact of doing that research presents a moderate chance of producing false alarms.

This raises a question of whether we are pursuing the correct model of value of information. It is certainly the correct one from the perspective of decision theory and practice. However, in that standard calculation, the expected value of the decisions given better information are compared to the "best" decision that the same decision maker would make if he or she were to make that decision without better information. This may be the theoretical ideal, and it is probably also the case for a corporate decision. However, with SRM one is concerned with a global public policy decision with many diffuse decision makers. The information in Tables 1 through 3 have been useful because they highlight how different people may be viewing the identical decision, by showing how values vary with individual, subjective judgments on probabilities. Optimists may be prepared to promote SRM as an important option for managing climate, but they may assign little value for the R&D associated with it, if the information resulting from the R&D is not close to perfect. Why might these same individuals therefore also argue in favor of doing that R&D? One possibility is that they realize that the R&D is the only possible way to find a path under which SRM might ever be accepted. They might understand that the greater pessimism about SRM among a large portion of the other individuals contributing to this policy decision will prevent SRM from being allowed unless and until the R&D is completed. Thus, the alternative expected value that the informed case would be compared to in order to calculate the value of information in a public policymaking setting might be the highly precautionary outcome of No SRM, rather than the optimized decision that uses a single set of priors. The precautionary attitude might also dominate the public policy decision if the R&D produces a "not safe" finding, such that SRM would never be allowed under that research outcome, with no consideration for the possibility that it could be a false positive. This would be a world in which those with relatively optimistic priors have to convince all of the others in society that SRM is a risk worth taking, and the value of information computed by an alternative model of choice such as this is very different than presented above.

Table 4 shows the estimates of the value of information that one would assign to research on the potential consequences of SRM, given one's own personal prior, but now assuming that the societal decision on SRM will be driven by a majority who are inclined to act in a precautionary manner if there is no concerted additional research to help them be persuaded otherwise, and who also will act in a precautionary manner if such research identifies a "not safe" signal. The settings of probabilities in Table 4 are identical to those in Table 3: it fixes the probability of a false positive at 25% and reports the alternative concept of value of information in a social choice setting for the full range of combinations of prior probabilities that SRM will be safe and probabilities of false negatives from the research.

Suddenly the large zone of zero value of information associated with relatively optimistic priors on SRM's consequences that was found in Table 3 is replaced by the highest value of information estimates we have estimated anywhere. In fact, these values are nearly as high as the value of SRM computed by Bickel and Lane. In a sense, in a societal decision making situation that may be dominated by precautionary attitudes, R&D may be the key that enables any consideration at all of undertaking a potentially highly valuable but also risky activity such as SRM. Thus, the R&D may have very high values to those who have strong preconceived views of the promise of a new and risky technology, even though traditional value of information theory would suggest that these same individuals ought to be the ones who would assign relatively low value of information to the same research.

Table 2. Value of Information Results in $ Trillions (Assuming Probability of False Positive, s = 0)

	0	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.05	0.8	0.6	0.4	0.2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.1	1.6					0.6	0.4	0.2	0	0	0	0	0	0	0	0	0	0	0	0
0.15	2.4							1		0.7	0.5	0.3	0.09	0	0	0	0	0	0	0
0.2	3.1									1.5			1	0.8		0.5	0.3	0.2	0.1	0.07
0.25	2.8					2	1.8				1.1		1	0.6		0.3	0.3			0.06
0.3	2.5													0.4	0.4	0.3				
0.35												0.5	0.4	0.4						
0.4	1.8						1	0.9		0.6	0.5	0.5				0.3				0.05
0.45							0.7	0.6	0.6	0.5										
0.5	1.1	1		0.8		0.6	0.6				0.4					0.2				0.04
0.55	0.8	0.7																		
0.6	0.7	0.7																		
0.65																				
0.7																				
0.75	0.4					0.3	0.3				0.2					0.1				0.02
0.8																				
0.85																				
0.9	0.2					0.1	0.1				0.08					0.04				
0.95	0.02					0.013	0.012				*	*	*	*	*	*	*	*	*	*
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Vertical axis is prior probability, p. Horizontal axis is Pr (False Negative) from R&D, q.

(*) implies values below $10b, which is the cost for the proposed R&D.

Color zones indicate the implied sequence of decisions through the decision tree.

- Green: never do SRM no matter what R&D conclusion is
- Blue: Plan early mitigation for no SRM, but do deploy SRM in 2025 if R&D concludes it will be safe.
- Yellow: Plan early mitigation for SRM, but deploy SRM only if R&D concludes it will be safe
- Orange: Do SRM even if R&D concludes it will not be safe (In Table 2, this zone actually only applies if prior p equals exactly 1.0. For values of p on [.95, 1.0), the decision is as defined by the yellow zone.)

Table 3. Value of Information Results in $ Trillions (Assuming Probability of False Positive, s = 0.25)

	0	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.05	0.6	0.4	0.2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.1	1.2			0.6	0.4	0.2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.15	1.8						0.6		0.3	0.07	0	0	0	0	0	0	0	0	0	0
0.2	2.3										0.6	0.4	0.2	0.02	0	0	0	0	0	0
0.25	1.8					1	0.8			0.3	0.1	0	0	0	0	0	0	0	0	0
0.3	1.3						0.3		0.04	0	0	0	0	0	0	0	0	0	0	0
0.35	0.8	0.6		0.3		0.03	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.4	0.2	0.09	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.45	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.55	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.65	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.75	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.85	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.95	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Vertical axis is prior probability, p. Horizontal axis if Pr (False Negative) from R&D, q.

Color zones indicate the implied sequence of decisions through the decision tree.

- Green: never do SRM no matter what R&D conclusion is
- Blue: Plan early mitigation for no SRM, but do deploy SRM in 2025 if R&D concludes it will be safe.
- Orange: Do SRM even if R&D concludes it will not be safe.

Table 4. Value of Information for Societal Decision in $ Trillions (Assuming Probability of False Positive, s = 0.25)

	0	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.05			0.2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.1	1.2					0.2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0.15										0.06	0	0	0	0	0	0	0	0	0	0
0.2														0.04	0	0	0	0	0	0
0.25	2.9					2.1					1.3					0.4	0.3	0.1	0	0
0.3																	1	0.9	0.8	0.8
0.35																	1.8	2.6		
0.4															2.8	2.7	2.6			
0.45														3.7	3.6	3.5				
0.5	5.9										4.8		4.6	4.4		4.4				4.4
0.55						6.5					5.5	5.4	5.3							
0.6						6.6				6.3	6.2									
0.65						7.3		7.1		7.1										
0.7	8.3			8.1	8	8														
0.75	8.9					8.8					8.6					8.6				
0.8																				
0.85																				
0.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9	10.9
0.95																				
1	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6	12.6

Vertical axis is prior probability, p. Horizontal axis if Pr (False Negative) from R&D, q.

Color zones indicate the implied sequence of decisions through the decision tree.

- Green: never do SRM no matter what R&D conclusion is
- Blue: Plan early mitigation for no SRM, but do deploy SRM in 2025 if R&D concludes it will be safe.
- Yellow: Plan early mitigation for SRM, but deploy SRM only if R&D concludes it will be safe

CONCLUSION

This Perspectives Paper has taken a first step to incorporate concerns about potential unintended consequences into the benefit and cost analysis of the Assessment Paper on geoengineering by Bickel and Lane, and to directly assess the value of an R&D program that would help reduce such uncertainties before any geoengineering deployment decision might be made. It has found that the value of ***perfect*** information would be much higher than their proposed research budget, but it also showed that ***imperfect*** information may have zero value. The value of information analysis thus would appear to undercut the assertion by the Assessment Paper authors that the value of SRM is so large that uncertainties are unlikely to affect the merits of conducting the further research needed to position SRM as a real option for potential deployment. This paper, however, has also proposed a possible alternative value of information calculation that may be more appropriate to use for societal decision making by groups of people who hold very different sets of priors. When the suggested alternative for calculating value of information in a social choice setting was applied, the value of the geoengineering R&D program exceeded even the value of perfect information over a very wide range of priors, and regions of zero value were confined to a narrow range consistent with a very pessimistic view of both the probability that the geoengineering will be safe and the probability that R&D will be able to identify such hazards in advance of full deployment.

This initial value of information analysis has addressed only one crucial uncertainty. It has hopefully created some interesting perspectives for discussion of the geoengineering Assessment Paper. It has also defined an analytical structure that can be expanded into a more complete form, as needed. The proposed R&D program would presumably reduce uncertainties over a variety of other uncertainties such as cost, functionality, climate change damages, etc., and thus value of information from this single program might be larger or smaller than the estimates here. If a more complete quantitative analysis is desired, the structure developed in this paper can serve as a foundation on which to build the more comprehensive analysis.

In my opinion, one of the most important next steps for this line of analysis would be to expand it to include AC as an additional option that could be deployed with or instead of SRM. That decision would be delayed until R&D is completed and has better characterized their respective true costs and risks. Because the uncertainties and risks of AC and SRM are largely independent of each other, both may contribute value, and in different ways than would be apparent from a comparison of their respective deterministic benefit-cost ratios. For example, AC might be found to have value as a backup option that could be deployed ***if*** SRM is used and then must be stopped suddenly.

It is also important to point out that the structuring of a decision problem is most valuable in its ability to get experts and decision makers to converse and bring new issues to the table when trying to communicate the decision problem in the highly structured format of a decision tree. That suggests that further work to expand on these foundations would probably best be done in a collaborative manner that draws in the comments, suggestions, and reactions of the range of experts, policy makers, and stakeholders engaged in the geoengineering issue.

ADDENDUM: BASIS FOR COST ESTIMATES USED IN FIGURE 3

Category	Summary of Estimates and Rationale
SRM Damages	$0 if full deployment of SRM is not done, or if it is done and is found to be safe. $1t if SRM is deployed and found to cause damages during it first ten years (i.e., 2025-2035). This is consistent with the example in Figure 2, but alternative assumptions can be explored later, as was done in Table 1.
SRM Cost	$0.7t in cases where SRM is deployed in 2025 and found to be safe, taken from Table 10 of the Assessment Paper. $0.3t in cases where SRM is deployed but stopped in 2035. (A more precise value could be provided by Bickel & Lane.) $0 if SRM never fully deployed. (Small-scale field tests of SRM prior to 2025 are not included, as these are considered to be part of the R&D budget whose value is being assessed in this analysis.)
GCC Damages	$3.6t in cases where SRM is deployed and found to be safe, and $11.9t in cases where SRM is never deployed, both taken from Table 4 of the Assessment Paper. In cases where SRM is deployed but then stopped in 2035, cost set to $15t on the assumption that the much more rapid changes in temperature that would occur post-2035 due to higher CO2 concentrations accumulated under the SRM strategy would exacerbate climate change damages relative to the case without any SRM. (A more precise value consistent with the result of the analysis assumptions could be obtained with another DICE run.)
Mitigation Cost	$0.6t when SRM is deployed and found safe and pre-2025 mitigation levels are tailored for this SRM outcome. $10.9t in cases where pre-2025 mitigation levels assume no future SRM benefits, and SRM is indeed not deployed. Both of the former values are taken from Table 4 of the Assessment Paper. In cases where pre-2025 mitigation assumes SRM will not occur, but SRM is deployed and found safe post-2025, it is assumed that 20% of the PV of the $10.9t cost will have been incurred by 2025 and this amount is added to the $0.6t associated with the SRM case, for an input value of $2.8t. (A more precise value could be estimated with another DICE run.) In cases where pre-2025 mitigation levels are tailored for an SRM deployment, but it is learned in 2020 that SRM will not be deployed, PV of mitigation costs is assumed to be about 25% higher than if the pre-2025 mitigation levels had been consistent with no SRM, due to the need to make up for lost time starting in 2020 to still avoid exceeding 2oC. This produces a cost of $13.6t. (A more precise estimate also could be obtained from another DICE run.) In the cases where SRM is deployed, but must be stopped in 2035, the cost is assumed to be $11t. This does not appear to be as high as in the former case because although both cases involve rapid increases in mitigation rates to make up for lost time, the former occurs 15 years earlier, and thus is less discounted in the present value. (A more precise estimate could be obtained from the same DICE run that could better inform the level of GCC damages for this case.)

COPENHAGEN CONSENSUS ON CLIMATE

The science is clear. Human-caused global warming is a problem that we must confront.

But which response to global warming will be best for the planet? The Copenhagen Consensus Center believes that it is vital to hold a global discussion on this topic.

The world turned to scientists to tell us about the problem of global warming. Now, we need to ensure that we have a solid scientific foundation when we choose global warming's solution. That is why the Copenhagen Consensus Center has commissioned research papers from specialist climate economists, outlining the costs and benefits of each way to respond to global warming.

It is the Copenhagen Consensus Center's view that the best solution to global warming will be the one that achieves the most 'good' for the lowest cost. To identify this solution and to further advance debate, the Copenhagen Consensus Center has assembled an Expert Panel of five world-class economists – including three recipients of the Nobel Prize –to deliberate on which solution to climate change would be most effective.

It is the Copenhagen Consensus Center's hope that this research will help provide a foundation for an informed debate about the best way to respond to this threat.

COPENHAGEN CONSENSUS CENTER

The Copenhagen Consensus Center is a global think-tank based in Denmark that publicizes the best ways for governments and philanthropists to spend aid and development money.

The Center commissions and conducts new research and analysis into competing spending priorities. In particular it focuses on the international community's efforts to solve the world's biggest challenges.

www.copenhagenconsensus.com