

Do probabilistic expert elicitations capture scientists' uncertainty about climate change?

Antony Millner · Raphael Calel · David
A. Stainforth · George MacKerron

Received: date / Accepted: date

Abstract Expert elicitation studies have become important barometers of scientific knowledge about future climate change (Morgan & Keith, 1995; Reilly et al., 2001; Morgan et al., 2006; Zickfeld et al., 2007; Kriegler et al., 2009; Zickfeld et al., 2010). Elicitations incorporate experts' understanding of known flaws in climate models, thus potentially providing a more comprehensive picture of uncertainty than model-driven methods. The goal of standard elicitation procedures is to determine experts' subjective probabilities for the values of key climate variables. These methods assume that experts' knowledge can be captured by subjective probabilities – however, foundational work in decision theory has demonstrated this need not be the case when their information is ambiguous (Ellsberg, 1961). We show that existing elicitation studies may qualitatively understate the extent of experts' uncertainty about climate change. We designed a choice experiment that allows us to empirically determine whether experts' knowledge about climate sensitivity (the equilibrium surface warming that results from a doubling of atmospheric CO₂ concentration) can be captured by subjective probabilities. Our results show that, even for this much studied and well understood quantity, a non-negligible proportion of climate scientists violate the choice axioms that must be satisfied for subjective probabilities to adequately describe their beliefs. Moreover, the cause of their violation of the axioms is the ambiguity in their knowledge. We expect these results to hold to a greater extent for less understood climate variables, calling into question the veracity of previous elicitations for these

A. Millner

Department of Agricultural and Resource Economics, University of California, Berkeley
Grantham Research Institute on Climate Change and the Environment, London School of
Economics and Political Science
Tel.: +44 207 107 5211
E-mail: A.Millner@lse.ac.uk

R. Calel, D.A. Stainforth, G. MacKerron

Grantham Research Institute on Climate Change and the Environment, London School of
Economics and Political Science

quantities. Our experimental design provides an instrument for detecting ambiguity, a valuable new source of information when linking climate science and climate policy which can help policy makers select decision tools appropriate to our true state of knowledge.

Keywords Expert Elicitation · Uncertainty · Ambiguity · Climate Sensitivity

Climate models and observational data are core elements of climate science, but predictions based on them are known to suffer deficiencies (Stainforth et al., 2007; Knutti, 2008) arising from the intrinsic difficulties in predicting high dimensional nonlinear systems (Smith, 2002, 2007) and incomplete understanding of relevant physical processes (Zickfeld et al., 2010). Careful expert elicitations can partially account for these difficulties, as experts are aware not only of models and their predictions, but also their relative strengths and weaknesses, and can factor this ‘meta-knowledge’ into their probabilistic assessments (Reilly et al., 2001). While human probability estimates are subject to biases (Kahneman et al., 1982), carefully conducted elicitations can minimize their impact on results (Morgan & Henrion, 1992).

The purpose of most elicitation studies is to attempt to estimate experts’ subjective probabilities for the values of key physical parameters, climate sensitivity being a much studied example. Subjective probabilities are assumed to completely capture ‘degrees of belief’ about the likely values of these parameters. The most satisfactory ontology of subjective probabilities is that offered by modern decision theory (Ramsey, 1931; De Finetti, 1937), in particular axiomatic subjective expected utility (SEU) theory (Savage, 1954), which forms the bedrock of a myriad of applications in the social sciences, including economic analysis of climate policy (Stern, 2007). In this scheme, experts reveal their subjective probabilities through choices over bets with uncertain outcomes; one works backwards from observed choices and principles of rational choice to infer subjective probabilities. Subjective probabilities are thus derivable from objective observations, placing them on a sound operational footing.

Not all beliefs, however, can be represented by subjective probabilities. In 1961 Daniel Ellsberg (Ellsberg, 1961) showed that many people (even eminent decision theorists) prefer to violate the axioms of SEU¹ in situations of deep uncertainty, where information about the likelihoods of alternative outcomes is incomplete, inconsistent, or non-existent. One of Ellsberg’s classic choice problems is described in Figure 1. Many people are not dissuaded from the choices depicted in the figure even after their violation of the axioms of SEU is pointed out to them (Ellsberg, 1961; Slovic & Tversky, 1974). A large literature (e.g. Gilboa et al. (2009); Binmore (2009); Gilboa (2009) and references therein)

¹ In fact, one of Ellsberg’s choice problems (see Figure 1 below) rules out preference representations much more general than SEU. The choices described in Figure 1 are inconsistent with any *probabilistically sophisticated* (Machina & Schmeidler, 1992) preference representation.

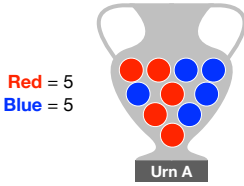

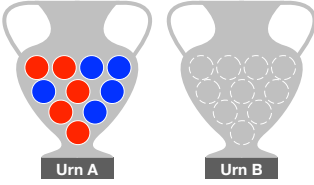
	Bets offered	Preference	What's implied?
 <p>Red = 5 Blue = 5</p> <p>Urn A</p>	<p>Win \$100 if Red is drawn</p> <p>or</p> <p>Win \$100 if Blue is drawn</p>	<p>Indifferent</p>	<p>$\Pr(\text{Red drawn from A}) = 0.5$</p>
 <p>Red = ? Blue = ? Red + Blue = 10</p> <p>Urn B</p>	<p>Win \$100 if Red is drawn</p> <p>or</p> <p>Win \$100 if Blue is drawn</p>	<p>Indifferent</p>	<p>$\Pr(\text{Red drawn from B}) = 0.5$</p>
 <p>Urn A Urn B</p>	<p>Win \$100 if Red is drawn from A</p> <p>or</p> <p>Win \$100 if Red is drawn from B</p>	<p>Prefer Red in A</p>	<p>$\Pr(\text{Red drawn from A}) > \Pr(\text{Red drawn from B})$</p>

Fig. 1 The Ellsberg Problem. Imagine that you are presented with Urn A, containing 10 balls, 5 red and 5 blue. A single ball will be drawn from the urn at random. A RED (BLUE) bet pays out \$100 if a red (blue) ball is drawn from the urn, and zero otherwise. Clearly, it is reasonable to be indifferent between the RED and BLUE bets. Since the payoffs in each of these bets are equal (but opposite), indifference between RED and BLUE reveals a subjective probability of 0.5 that the ball will be red. Now consider Urn B, again containing 10 balls, each either red or blue, however this time in *unknown proportions*. When faced with a choice between RED or BLUE bets in this urn, many people again express indifference, allowing us to again infer a 0.5 subjective probability that a ball drawn from this urn will be red. Finally, when asked whether they would prefer a RED bet on Urn A, or a RED bet on Urn B, many people express a strict preference for betting on Urn A, preferring to bet on a known rather than an unknown risk. However, according to the tenets of subjective probability, these bets are identical, since their payoffs are the same, and it has been established that the subjective probability of drawing a red ball is the same in each urn. We thus arrive at a contradiction – there is no set of subjective probabilities that can describe both indifference between RED and BLUE on Urns A and B, and a strict preference for RED in Urn A over RED in Urn B.

has argued that such choice behaviour, referred to as ambiguity aversion, is rational in situations of informational paucity.

It is thus an empirical matter to determine whether subjective probabilities accurately capture experts' knowledge. If experts exhibit ambiguity aversion over bets on the quantity of interest we qualitatively understate the extent of their uncertainty if we force their beliefs into the straitjacket of subjective probabilities. Existing elicitations almost invariably presuppose the existence

of subjective probabilities², since experts' cumulative distribution functions (CDF) are elicited directly, instead of inferred from choices over bets with uncertain outcomes. Such choice tasks can allow experts to express a richer variety of beliefs, and can be used as diagnostics to determine the quality of their information. If experts' knowledge is indeed ambiguous, this is important, policy relevant, information – the decision tools appropriate for policy evaluation under ambiguity are substantially different from those that apply when subjective probabilities are available (Ellsberg, 2001; Binmore, 2009; Gilboa, 2009).

The Ellsberg Problem in Figure 1 can be modified to detect ambiguity in experts' scientific knowledge (Heath & Tversky, 1991). We focussed on climate sensitivity³ (S), as it is possibly the most familiar parameter of climate change science, is widely studied in both the expert elicitation and conventional scientific literature, and is much used in economic analysis of climate policy (Stern, 2007; Nordhaus, 2008).

Values of S corresponding to the 5th, 50th, and 95th percentiles of each scientist's hypothesized CDF for S were initially elicited using standard direct probabilistic elicitation techniques (Morgan & Henrion, 1992). The results of this initial portion of the survey are presented in Figure 2, alongside results from two previous such elicitations, and a collection of model-based probability distributions for climate sensitivity taken from Meinshausen et al. (2009). The figure shows that our results are similar to those obtained in the recent study of Zickfeld et al. (2010)⁴. Both our study and the Zickfeld et al. (2010) study show a substantial increase in estimates of the percentiles relative to the early study of Morgan & Keith (1995).

In the second half of our survey the experts completed four sets of betting tasks – three sets of bets on the value of S (one set at each of the expert's elicited percentiles), and one set of bets on the Ellsberg Problem. We will refer to the three betting tasks on the value of S as the Climate Problem. For each task in the Climate Problem, experts were asked to make three choices: first between bets on the colour of a ball drawn from an urn of known composition (calibrated to match the percentile of S), then between comparable bets on the value of S , and finally between a bet on the urn with known composition and a bet on the value of S . In order to clarify this part of the experiment, consider the following simple example:

Suppose that an expert's estimate for the median of the distribution for S is $S_{0.5}$. We presented this expert with the following betting tasks:

² Kriegler et al. (2009) is an exception, however its results are difficult to interpret since it prompted experts for a range of probabilities (taking the existence of imprecise probabilities for granted), instead of inferring the non-existence of subjective probabilities from observed choices.

³ The precise definition of climate sensitivity we used is quoted in the Supplementary Information, and is also available on the survey website.

⁴ We computed a two-sided Wilcoxon rank-sum test at each of the three percentiles to test the hypotheses that the percentile estimates in our study and Zickfeld et al. (2010) were drawn from the same sampling distributions. All of the P-values from the three tests exceed the Bonferroni corrected 5% threshold.

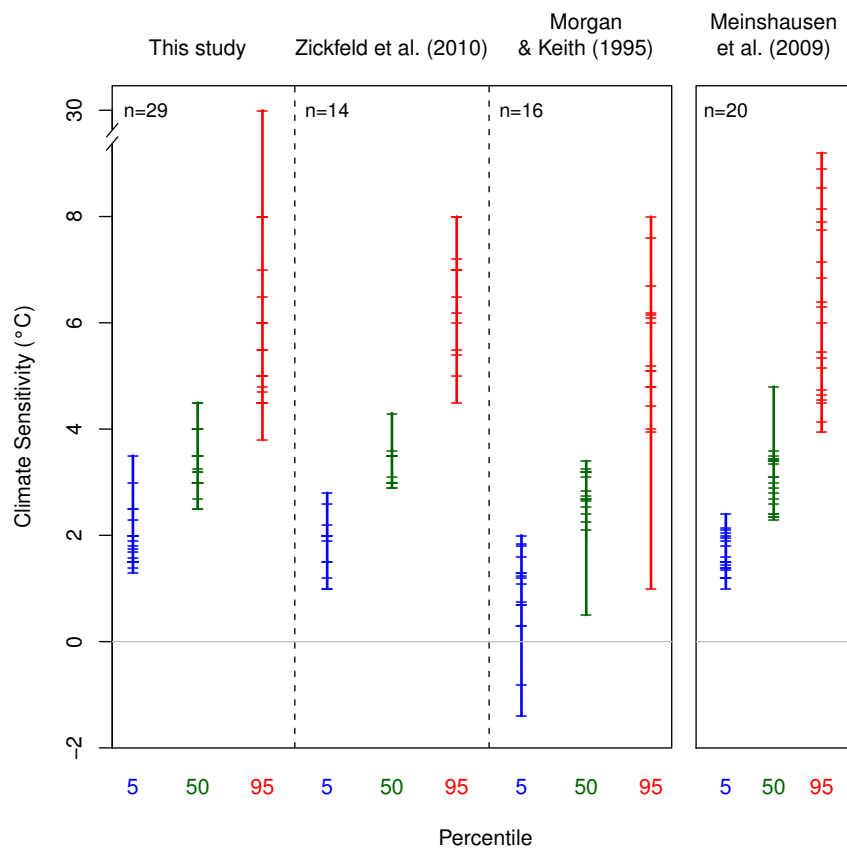


Fig. 2 Comparison of elicited percentiles of the probability distribution for climate sensitivity in this study with two previous elicitations (Zickfeld et al., 2010; Morgan & Keith, 1995), and with the model based estimates collated by Meinshausen et al. (2009). The number of tick marks at each percentile in the elicitation studies does not coincide with the size of the sample, as several experts reported overlapping values.

Choice I: Urn I contains 100 balls, 50 red and 50 blue. Choose between bets A and B , or are you indifferent between them?

- A : Win \$50 if a red ball is drawn, and zero otherwise.
- B : Win \$50 if a blue ball is drawn, and zero otherwise.

Choice II: Consider the true value of climate sensitivity S . Choose between bets C and D , or are you indifferent between them?

- C : Win \$50 if $S > S_{0.5}$, and zero otherwise.
- D : Win \$50 if $S < S_{0.5}$, and zero otherwise.

Suppose that the expert was indifferent between A and B , and also indifferent between C and D . We would then present her with a third betting task:

Choice III: Consider the true value of climate sensitivity S , and Urn I. Choose between bets A and C , or are you indifferent between them?

- A : Win \$50 if a red ball is drawn from Urn I, and zero otherwise.
- C : Win \$50 if $S > S_{0.5}$, and zero otherwise.

In exact analogy with the Ellsberg Problem represented in Figure 1, if the expert exhibits a strict preference for A over C , or indeed C over A , she has violated SEU, and her preferences are inconsistent with the existence of subjective probabilities. In this manner, a set of three choices is presented to each expert at each of her elicited percentiles. The construction of the bets and inference of SEU violations at the 5th and 95th percentiles is more complex than in this simple example – interested readers may consult the Supplementary Information for further details. After the experts had completed the Climate Problem, they also completed a set of betting tasks on the classic Ellsberg Problem of Figure 1. The results of both these sets of betting tasks – the Climate Problem and the Ellsberg Problem – are summarized in Figure 3.

To interpret the results in Figure 3, begin by noting that violation of SEU in the Ellsberg Problem indicates that an expert prefers to bet on known rather than unknown risks – she is ambiguity averse. Ambiguity aversion on the Ellsberg Problem provides a potential causal explanation for violations of SEU on the Climate Problem. That is to say, SEU violations in the Ellsberg Problem reveal experts’ *attitudes* towards deep uncertainty. If we then observe similar violations of SEU on the Climate Problem, we can use what we know about their attitudes to deep uncertainty (as revealed in the Ellsberg Problem) to try to understand whether these violations of SEU are due to the presence of ambiguity about S . In order to establish that ambiguity is indeed the cause of SEU violations for bets on S , we need to perform statistical tests to determine whether there is a dependence between SEU violations on the Ellsberg Problem and SEU violations on the Climate Problem. Thus our analysis does not assume that any single SEU violation on the Climate Problem, which may be due to idiosyncratic factors, is definitive evidence of ambiguity about S . Rather, we treat SEU violations as a noisy signal, and use statistical analysis to infer the presence of ambiguous knowledge about S in the sample. To examine these issues quantitatively it is useful to summarize the results of the experiment in a contingency table (Table 1).

17% of the sample violated SEU on bets on S , while 31% of the sample violated SEU on the Ellsberg Problem⁵ ($N = 29$). From Table 1 we see that

⁵ Experts in our sample were less likely to violate SEU on the Ellsberg Problem than in most other published studies, where SEU violation rates can be up to 80% (Slovic & Tver-

Table 1 Contingency table for SEU violations on the Ellsberg Problem (EP) and Climate Problems (CP) ($N = 29$)

	Violation of SEU on CP	No violation of SEU on CP
Violation of SEU on EP	3	6
No violation of SEU on EP	2	18

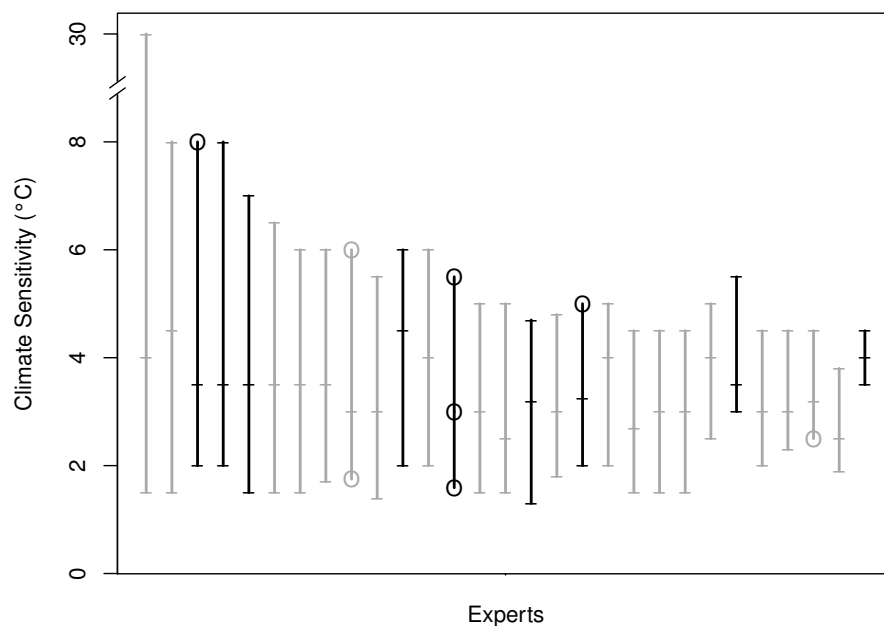


Fig. 3 Summary of results from the Climate and Ellsberg Problems. Dark lines are those experts who exhibited ambiguity aversion on the Ellsberg Problem described in Figure 1, while light lines are those who did not. Tick marks on each line indicate the 5th, 50th, and 95th percentile of the expert's CDF for climate sensitivity, as elicited by conventional direct methods. Circles indicate that the expert violated SEU on the Climate Problem at the corresponding percentile of S .

33% of the experts who violated SEU on the Ellsberg problem also violated SEU on at least one bet⁶ in the Climate Problem ($N = 9$), while 10% of the

sky, 1974; Camerer & Weber, 1992). This most likely reflects the scientist's mathematical training, and suggests that those SEU violations that we do observe are likely not due to lack of familiarity with the rules of probability theory.

⁶ Note that all we need to show is that SEU is violated once in order to conclude that an expert's knowledge cannot be described by subjective probabilities. One might argue that it is easy to observe at least one violation of SEU by simply asking experts to make bets on a large number of values of S , however the larger the number of bets, the less power the experimental design has to detect correlations between behavior in the Ellsberg Problem and that in the Climate Problem. The current design achieves a balance between detecting SEU violations, and preserving sufficient statistical power to allow us to ascribe them to the presence of ambiguity.

experts who did not violate SEU on the Ellsberg Problem violated SEU on at least one bet in the Climate Problem ($N = 20$). If SEU violations in the Climate Problem were *not* caused by ambiguity about S , the proportion of SEU violators on S would not depend on whether experts violated SEU on the Ellsberg Problem. Thus our null hypothesis is that the proportion of SEU violators on the Climate problem should be the same for SEU violators on the Ellsberg Problem and for non-SEU violators on the Ellsberg Problem. Using Barnard’s test (Barnard, 1945) – an exact test of independence valid for small samples, and with high power for 2×2 contingency tables – we reject this null hypothesis at the 10% level (one-sided test⁷, $P = 0.083$). This suggests that it is indeed the ambiguity in S that drives SEU violations in the Climate Problem. Theory predicts that the upper tail of the distribution for S is less constrained by the instrumental record than the rest of the distribution (Allen et al., 2006; Roe & Baker, 2007), and is thus more likely to be ambiguous. Our results are also consistent with this prediction – the strength of the association between SEU violations on the Ellsberg Problem and SEU violations on S is increasing in the percentiles of S (see Supplementary Information). If we restrict the analysis only to SEU violations on bets at the 95th percentile for S the results are significant at the 5% level (one-sided test, $P = 0.045$). Finally, the data also illustrate the difference between risk as captured by the width of the elicited subjective probability distribution, and ambiguity. We cannot reject the hypothesis that SEU violations in the Climate Problem are independent of whether experts’ elicited 5-95% range for S was above or below the sample median ($P = 0.574$, see Supplementary Information).

We thus observe three patterns in the data:

1. There is a positive dependence between SEU violations in the Ellsberg Problem and SEU violations in the Climate Problem.
2. The evidence for dependence between SEU violations on the Ellsberg Problem and SEU violations on the Climate Problem is strongest in the upper tail of the elicited distributions for S .
3. SEU violations in the Climate Problem are not significantly dependent on ‘risk’, as measured by the width of the elicited distribution for S .

The hypothesis that those experts who violate SEU on the Climate Problem are doing so because their knowledge of the distribution of S is ambiguous is consistent with all these patterns. Although our results reflect the opinions of a moderate number of experts, our statistical analysis is exact, and does not rely on asymptotic methods that only hold for large samples.

Our experimental design was intentionally conservative, focussing on one of the most studied quantities in climate science, and is likely to understate the presence of ambiguity (see Supplemental Information for a discussion). It

⁷ We use a one-sided test as our hypothesis is that ambiguous beliefs about climate sensitivity cause SEU violations on the Climate Problem to be *more* likely amongst those who violate SEU on the Ellsberg Problem than amongst those who do not. Thus our alternative hypothesis is ‘positive dependence’ between SEU violations on the Ellsberg and Climate Problems.

is likely that elicitation conducted for more geographically localized, or less understood, climate variables (e.g. transient climate response (Zickfeld et al., 2010), aerosol forcing (Morgan et al., 2006), Atlantic meridional overturning (Zickfeld et al., 2007)) are more susceptible to the effects of ambiguity. Nevertheless, we believe that expert elicitation studies do contain useful information, and view our methods as a complementary test of the quality of elicited probabilities which allows us to detect ‘uncertainty about uncertainty’. Detecting ambiguity in experts’ beliefs does not imply that their opinions cannot be captured by elicitation, but rather that alternative methods, e.g. imprecise probability estimates (Walley, 1990; Kriegler et al., 2009), may be more appropriate. We view tests such as ours as a precondition for applying such methods however – one must reject the existence of subjective probabilities before prompting experts for imprecise probabilities.

Moreover, the existence of ambiguity does not imply ignorance, and should not be seen as an excuse for inaction. Decision theory provides several tools for policy analysis when knowledge is ambiguous (Savage, 1954; Resnik, 1987; Gilboa, 2009; Binmore, 2009). If anything, it is likely that accounting for the ambiguity in our knowledge recommends stronger mitigation policies than those based on conventional probabilistic decision tools (Millner et al., 2010; Lemoine & Traeger, 2012; Woodward & Bishop, 1997). Our method is of course not restricted to applications in climate science - it applies to any elicitation exercise and could be valuable in understanding how to use scientific information in a variety of policy situations. Including tests for the presence of ambiguity in future elicitation would provide more faithful representations of experts’ knowledge and enable us to select policies that are more consistent with our knowledge of, and attitude towards, uncertainty.

A Methods Summary

Participants were recruited by e-mail and word of mouth over a period of 3 months beginning in December 2010. 42 respondents completed the survey (available in full online at <http://www.climate.websperiment.org>). All respondents consented to be identified as participants, and were informed that their responses would be anonymized. 13 respondents were removed from the sample as they either stated that they are not familiar with the literature on climate sensitivity estimation, or were not primarily engaged in climate science research at the time of the survey.

The 29 experts in our sample were: Gab Abramowitz, James Annan, Kyle Armour, David Easterling, Seita Emori, John Fasullo, Chris Folland, Chris Forest, Piers Forster, John Harte, Gabriele Hegerl, Gregory Jones, Reto Knutti, Gerald Meehl, James Murphy, Falk Niehoerster, Geert Jan van Oldenborgh, John Reilly, Gerard Roe, Ben Sanderson, Stephen Schwartz, Carolyn Snyder, Andrei Sokolov, Claudia Tebaldi, Simon Tett, Warren Washington, Andrew Weaver, Rob Wilby, Carl Wunsch. All reported results have been anonymized.

Each experts’ hypothesized 5th, 50th and 95th percentile of the distribution for S were initially elicited using standard probabilistic elicitation methods. They then completed four sets of betting tasks – three on the Climate Problem (one at each of the elicited percentiles of S), and one on the Ellsberg Problem. Full details of these betting tasks are available in the Supplementary Information. Participants could move back and forth through the survey at any time, and had access to help boxes on each screen with reminders about quantity

definitions and judgmental biases to be aware of when forming their answers. They could also change their answers at any time. We used data only from those experts who completed the survey in full. There was no time limit on the survey, and experts were informed that they should take as much time as they need to form their best judgments. The Ellsberg Problem was presented at the very end of the survey, so as not to prime participants to think in terms of ambiguity.

B Author contributions

AM conceived of the research. RC and AM designed the experiment. GM implemented the online survey. DAS provided guidance on the formulation of the survey questions. RC, AM, and DAS recruited participants and ran the experiment. RC analyzed the data, and AM wrote the paper.

Acknowledgements AM was supported by a Ciriacy-Wantrup postdoctoral fellowship at UC Berkeley during the course of this work. RC is supported by the UK Economic and Social Research Council (ESRC) and the Jan Wallander and Tom Hedelius Foundation. DAS acknowledges the support of the LSE’s Grantham Research Institute on Climate Change and the Environment and the ESRC Centre for Climate Change Economics and Policy, funded by the Economic and Social Research Council and Munich Re. We thank Rachel Denison for advice and comments.

References

- M. Allen, et al. (2006). ‘Observational constraints on climate sensitivity’. In H. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley, & G. Yohe (eds.), *Avoiding dangerous climate change*, p. 406. Cambridge University Press, Cambridge, UK.
- G. A. Barnard (1945). ‘A New Test for 2×2 Tables’. *Nature* **156**(3954):177–177.
- K. Binmore (2009). *Rational Decisions*. Princeton University Press.
- C. Camerer & M. Weber (1992). ‘Recent developments in modeling preferences: Uncertainty and ambiguity’. *Journal of Risk and Uncertainty* **5**(4):325–370.
- B. De Finetti (1937). *La prévision : ses lois logiques, ses sources subjectives*. Annales de l’Institut Henri Poincaré. Paris: Institut Henri Poincaré.
- D. Ellsberg (1961). ‘Risk, Ambiguity, and the Savage Axioms’. *The Quarterly Journal of Economics* **75**(4):643–669.
- D. Ellsberg (2001). *Risk, Ambiguity and Decision*. Routledge.
- I. Gilboa (2009). *Theory of Decision under Uncertainty*. Cambridge University Press, 1 edn.
- I. Gilboa, et al. (2009). ‘Is It Always Rational to Satisfy Savage’s Axioms?’. *Economics and Philosophy* **25**(3):285–296.
- C. Heath & A. Tversky (1991). ‘Preference and belief: Ambiguity and competence in choice under uncertainty’. *Journal of Risk and Uncertainty* **4**(1):5–28.
- D. Kahneman, et al. (1982). *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- R. Knutti (2008). ‘Should we believe model predictions of future climate change?’. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **366**(1885):4647–4664.
- E. Kriegler, et al. (2009). ‘Imprecise probability assessment of tipping points in the climate system’. *Proceedings of the National Academy of Sciences* **106**(13):5041–5046.
- D. Lemoine & C. P. Traeger (2012). ‘Tipping points and ambiguity in the Economics of Climate Change’. *NBER working paper no. 18230*.
- M. J. Machina & D. Schmeidler (1992). ‘A More Robust Definition of Subjective Probability’. *Econometrica* **60**(4):745–780.

- M. Meinshausen, et al. (2009). 'Greenhouse-gas emission targets for limiting global warming to 2C'. *Nature* **458**(7242):1158–1162.
- A. Millner, et al. (2010). 'Ambiguity and climate policy'. *NBER working paper no. 16050* (<http://www.nber.org/papers/w16050.pdf>).
- M. G. Morgan, et al. (2006). 'Elicitation of Expert Judgments of Aerosol Forcing'. *Climatic Change* **75**(1-2):195–214.
- M. G. Morgan & M. Henrion (1992). *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press.
- M. G. Morgan & D. W. Keith (1995). 'Subjective judgements by Climate Experts'. *Environmental Science & Technology* **29**(10).
- W. D. Nordhaus (2008). *A question of balance*. Yale University Press.
- F. Ramsey (1931). 'Truth and Probability'. In R. Braithwaite (ed.), *The foundations of mathematics and other logical essays*. London: Kegan, Paul, Trench, Trubner & Co., New York: Harcourt, Brace and Company.
- J. Reilly, et al. (2001). 'Uncertainty and Climate Change Assessments'. *Science* **293**(5529):430–433.
- M. D. Resnik (1987). *Choices: an introduction to decision theory*. U of Minnesota Press.
- G. H. Roe & M. B. Baker (2007). 'Why Is Climate Sensitivity So Unpredictable?'. *Science* **318**(5850):629–632.
- L. J. Savage (1954). *The Foundations of Statistics*. Wiley and Sons.
- P. Slovic & A. Tversky (1974). 'Who accepts Savage's axiom?'. *Behavioral Science* **19**(6):368–373.
- L. A. Smith (2002). 'What might we learn from climate forecasts?'. *Proceedings of the National Academy of Sciences of the United States of America* **99**(Suppl 1):2487–2492.
- L. A. Smith (2007). *Chaos : a very short introduction*, vol. 159. Oxford University Press, Oxford.
- D. Stainforth, et al. (2007). 'Confidence, uncertainty and decision-support relevance in climate predictions'. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365**(1857):2145–2161.
- N. H. Stern (2007). *The economics of climate change : the Stern review*. Cambridge University Press, Cambridge.
- P. Walley (1990). *Statistical Reasoning with Imprecise Probabilities*. Monographs on Statistics and Applied Probability 42. Chapman & Hall.
- R. T. Woodward & R. C. Bishop (1997). 'How to Decide When Experts Disagree: Uncertainty-Based Choice Rules in Environmental Policy'. *Land Economics* **73**(4):492–507.
- K. Zickfeld, et al. (2007). 'Expert judgements on the response of the Atlantic meridional overturning circulation to climate change'. *Climatic Change* **82**(3-4):235–265.
- K. Zickfeld, et al. (2010). 'Expert judgments about transient climate response to alternative future trajectories of radiative forcing'. *Proceedings of the National Academy of Sciences*