Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties

Frances C. Moore,[1*] Moritz A. Drupp,[2] James Rising,[3]
Simon Dietz,[4] Ivan Rudik,[5] Gernot Wagner[6]

[1]Department of Environmental Science and Policy, University of California Davis
[4] Department of Economics and Center for Earth System Research and Sustainability (CEN), University of Hamburg
[3] School of Marine Science and Policy, University of Delaware
[4] London School of Economics
[5] Dyson School of Applied Economics and Management, Cornell University
[6] Columbia Business School

*To whom correspondence should be addressed; E-mail: fmoore@ucdavis.edu.

Estimating the cost from a ton of CO_2 to society requires connecting a model of the climate system with representation of the economic and social effect of changes in climate and the valuation and aggregation of diverse, uncertain impacts across both time and space. The literature on this cost, termed the social cost of carbon (SCC) is large and growing, with substantial differences in underlying assumptions both across and within studies. Significant prior work has focused on better constraining parameter values such as climate sensitivity, the discount rate, and the damage function. A growing literature, however, has also examined the effect of varying more fundamental structural elements of the models supporting SCC calculations. These structural model choices - including the introduction of climate or economic tipping points, changing the structure of economic preferences, and the persistence of climate damages - have often been analyzed in piecemeal, uncoordinated fashion, leaving their relative importance unclear. Here we perform a comprehensive synthesis of the evidence on the SCC, combining 1823 estimates of the SCC from 147 studies published between 2000 and 2020 with a survey of the authors of these studies. We find that the distribution in published and expert SCCs are both wide and substantially right-tailed. Survey evidence suggests that experts believe there to be a substantial downward bias in published SCC values. Analysis of the drivers of variance in the distribution reveals that structural variation across SCC models is important, particularly the persistence of climate damages. We estimate a random forest model based on SCC variation in the literature and combine this with expert assessment to generate a 'synthetic SCC' distribution integrating over expert assessments of uncertainty in model structures and the discount rate as well as parametric and residual uncertainty represented in the literature. Thisdistributionhas a mean of \$467 per tCO_2 for a 2020 pulse year (5%–95% range: \$14–\$1379). There is thus a substantial and varied body of evidence pointing towards using a high SCC in policy-making.

1 Introduction

Anthropogenic climate change will affect the welfare of people around the world for centuries into the future. Because these costs are largely not Incorporated into energy or land-use decisions, climate change has been termed "the greatest and widest-ranging market failure ever seen" (*1, p. i*). Incorporating climate costs into the prices of economic activities that emit greenhouse gases, either through carbon pricing or other regulations, is essential for averting the worst climate outcomes. Quantifying these costs is extremely challenging as it involves projecting and valuing the effects of climate change in all countries and sectors far into the future, an exercise that is rife with uncertainties and contestation.

The external costs of carbon dioxide emissions are summarized by the 'social cost of carbon' (SCC): the present value of all future impacts from an additional ton of CO_2 emissions. The SCC is key for understanding the benefits of emissions-reduction policies and is used for climate and energy policy analysis in the United States, Germany, Canada, and several sub-national jurisdictions (*2, 3*). Integrated assessment models (IAMs) commonly used to calculate the SCC have been criticized on various grounds, including inaccurate climate and carbon-cycle modeling, ignoring irreversibilities and tipping points in the climate system, failing to adequately model uncertainty or the potential for catastrophic outcomes, and relying on dated science in the representation of climate impacts (*4–8*).

The continuing importance of the SCC as a tool for climate policy analysis (*2*) and recognition of failings in IAMs currently used to calculate it has led to a surge of research seeking to improve, expand, and update the estimates. Major strands of this literature include: improving modeling of Earth system dynamics (*9–12*); disentangling preferences over risk and time using more complex utility functions (*13–15*); representing tipping points and associated uncertainties in damages (*16–19*); addressing model uncertainty, ambiguity, and learning new information (*20–24*); allowing climate damages to affect the growth rate rather than just the level of economic output (*11, 25–27*); calibrating aggregate climate damages on recent economic and scientific evidence (*11, 20, 25, 28, 29*); modeling the distribution of climate damages and incorporating spatial inequality aversion (*30–32*); and allowing for climate damages to non-market goods, such as natural systems or cultural heritage, which are imperfectly substitutable with market-traded goods (*33–36*).

Although this literature is now substantial, it has accumulated piecemeal. The vast majority of papers make one or two structural model choices to a simpler IAM and report how these alter SCC values, often with an exploration of associated parametric uncertainty. The implications of the full suite of issues addressed by this literature have not been assessed, including what are the most significant drivers of the variance in SCC estimates. Previous syntheses have quantified the distribution of SCC estimates and explored a limited set of covariates, such as publication year and discounting (*37, 38*), as well as the possible role of publication bias (*39*). Previous modeling studies have made multiple simultaneous changes to individual IAMs (*12, 40*), or have undertaken systematic IAM inter-comparisons and evaluations (*41, 42*), albeit focusing on a limited number of IAMs with comparable model structures. Previous expert surveys have either imposed very specific structure or none at all (*43–45*), or have focused on carbon prices (*46*). Thus, prior studies only illuminate the role of a subset of mechanisms and structural model choices.

This paper provides the most comprehensive assessment to date of SCC estimates, including how structural modeling choices shape the SCC. It builds on two complementary approaches. First, we perform a systematic analysis of SCC values published between 2000 and 2020. After reviewing over 2800 abstracts, we identified 1823 estimates (or distributions of estimates) published in 147 studies. SCC estimates and, where reported, the distribution of parametric uncertainty were recorded, along with over 35 covariates capturing details of the estimate itself (e.g., SCC year, discounting scheme, and socio-economic and emissions scenarios), major structural model choices to the standard SCC modeling approach (e.g., growth rate damages, inequality aversion, and disentangled preferences over risk and time), and sources of parametric variation (e.g., distributions over productivity growth, climate sensitivity, and damage function parameters). Second, we conduct an expert survey of the authors of the SCC papers in our analysis. We elicit estimates of what experts perceive is the distribution of published SCC values in the literature, and the distribution of what they believe best approximates the "true" SCC, all things considered. We also ask experts to break down the wedge between their central estimates of the literature SCC and the "true" SCC, which allows us to investigate divergence between the distributions and their determinants. Furthermore, we elicit experts' views on how much a selection of structural modeling changes in the literature affects SCC estimates and their assessment of whether these model modifications *improve* SCC estimates.

Our study therefore contains two complementary data-generation processes: a meta-analysis, which collects much richer data on published SCC estimates and their determinants, and an expert survey. Their combination allows us to present the most comprehensive characterization yet of SCC estimates and their drivers. We furthermore combine these two lines of evidence using a random forest model to create a "synthetic SCC" that reflects both the SCC variation in the reported literature and expert assessments of discounting and structural model choices. Additional details on the literature review, coding of values, data cleaning and processing, and expert survey are provided in the Supplementary Information.

2 The SCC Distribution

The systematic review of the literature yields 1823 SCC estimates (or distributions) from 147 studies. Many studies report multiple SCC estimates. For each of the 1823 estimates, we collect information on the central SCC estimate, emission pulse year, discounting, damage function, economic and emissions scenario, structural assumptions, and distribution resulting from parametric uncertainty (where reported, 832 of the 1823 estimates).

To characterize the distribution of SCC values appearing in the published literature, we sample from the dataset using a hierarchical sampling scheme. We draw 10 million SCC values sampling uniformly from the 147 studies in the dataset, then uniformly from the set of estimates within each paper (i.e. unique SCC year-discounting-scenario-model structure combinations), and finally from the parametric uncertainty of each estimate, if applicable. Alternate sampling schemes that account for non-independence between papers using sets of shared authors, or for different quality of studies using a normalized citation-based weighting, give quantitatively similar distributions (Table S4).

Figure 1 gives the distribution of SCCs reported in the literature for pulse years between 2010 and 2030, which we use as the 2020 SCC equivalent sample. The variation in SCC values from the literature is substantial and asymmetric, exhibiting a long right tail, and a mean value (\$151 per tCO_2 after truncating the upper and lower 0.1% of values) that is several times higher than the median (\$41). Statistical tests show evidence for a heavy tail in the SCC distribution, echoing (*47*), with slope of the mean excess function greater than 1 and α values between 1 and 2, indicative of a distribution with infinite variance but finite mean (see Table S5 in the SI). Table S6 shows evidence that the upper tail of SCC values is particularly associated with papers integrating persistent or growth rate damages and certain types of parametric variation (specifically around adaptation rates, the damage function, and income elasticity).

Figure 1 also shows SCC distributions arising from two other lines of evidence: our survey of SCC authors conducted as part of this study and two distributions arising from U.S. government exercises. Survey participants were asked to provide both their estimates of the distribution of published SCC values (literature estimate, Figure 1), and their best estimate of a "true" or "comprehensive" SCC distribution – accounting for anything missing from the published literature or any systematic biases therein. These two distributions show that experts believe there is a substantial bias in the published literature, with the median estimate increasing from \$34 per tCO_2 for the literature estimate to \$85 per tCO_2 for the comprehensive estimate. Similarly to values from the literature, expert assessment shows a long right-tail on the comprehensive SCC distribution, with a mean of \$160 and a 95th percentile over \$580. Notably, while the mean of the literature distribution (\$151) and the expert comprehensive assessment ($160) are similar, the shapes of the distribution are quite distinct with the expert assessment putting substanatially higher probability on SCC values over $100 per tCO_2.

Figure 1 also shows the distribution from two U.S. government SCC estimates: the current distribution from the Inter-Agency Working Group on the Social Cost of Carbon, which uses the DICE, PAGE and FUND models and 2.5, 3, and 5% discount rates; this underpins the still formal US-SCC of around $50 per tCO_2; and a new proposed distribution from EPA which uses three updated damage estimates and 1.5, 2, and 2.5% discount rates, as well as updated distributions over socio-economic and emissions trajectories and climate system uncertainty (*3*). Figure 1 makes clear that multiple lines of evidence demonstrate that the current IWG SCC distribution is down-ward biased. For instance, the literature distribution, expert comprehensive assessment and more recent EPA analysis put the probability of an SCC over $100 per tCO_2 at 24.5%, 43.5% and 67.8% respectively, compared to only 9.4% in the current IWG estimate.

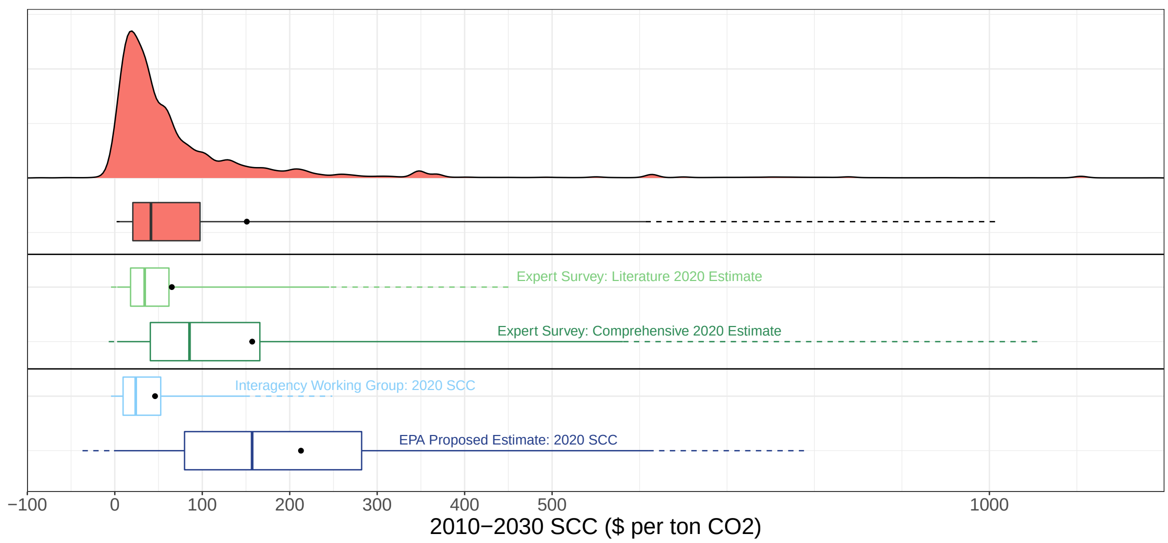

Figure 1: SCC distribution from the published literature (2020 $ per ton CO_2) for the 2010-2030 time period (pink). Central panel shows two distributions from the expert survey: respondents' estimate of the distribution of 2020 SCC values in the literature (light green) and their comprehensive assessment of the actual 2020 SCC distribution (dark green). Lower panel shows two distributions of U.S. Federal Government 2020 SCC estimates: the 2021 Interagency Working Group (light blue) and a proposed update by the U.S. Environmental Protection Agency (dark blue). Boxplots show the median (line), inter-quartile range (box), 5-95% range (solid line) and 2.5-97.5% range (dashed lines). Dots show the mean after trimming the upper and lower 0.1% of each distribution

2.1 SCC Distribution Under Alternate Structural Models

Figure 2a shows the 2010-2030 SCC distribution differentiated by nine major classes of structural models for SCC calculations found in the literature. (More detailed descriptions of these changes, along with example references integrating them are given in SI section S3). Other than papers incorporating learning, these changes tend to increase SCC values; largest effects are seen for estimates distinguishing between risk and time preferences via Epstein-Zin utility, estimates that allow for limited substitutability between market and non-market goods, and estimates allowing for persistent damages typically via impacts to economic growth. These raise the median SCC from \$38 per ton CO_2 in the reference distribution to \$57, \$128, and \$180 per ton CO_2 respectively. Effects on the upper tail of the distribution are also substantial; the 75th percentile of the distribution increases from \$70 in the reference distribution to \$108, \$197, and \$534 respectively. Table S12 shows evidence that these differences are not driven by systematic differences in discount rate or pulse year.

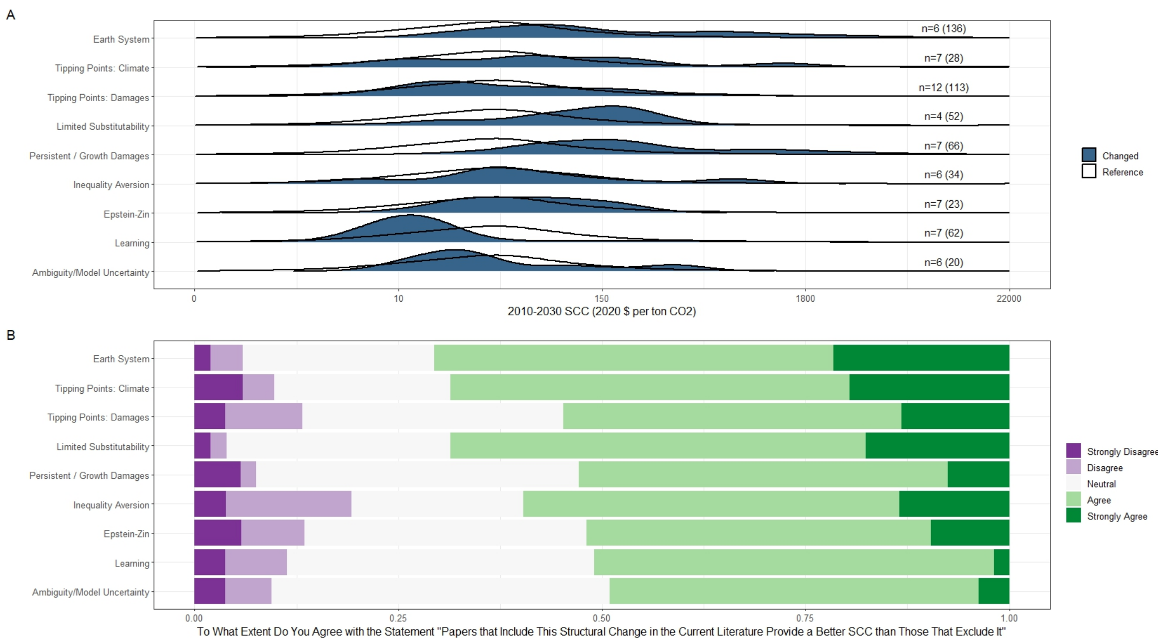

Figure 2: **The role of structural model choices in the literature SCC distribution.** a) SCC distribution for pulse years between 2010 and 2030 with (changed, blue) and without (reference, white) particular structural model choices to the SCC modeling. The reference distribution shows estimates without any of the 9 structural model choices and is the same in all rows. Numbers to the right of the graph give the number of papers making that structural model choices and, in parentheses, the total number of SCC observations making up the changed distribution in that row. Distributions are shown on a log scale, which requires dropping 2.9% of the distribution at or below \$0 per ton CO_2. b) Distribution of views from the expert survey asking respondents for their assessments of the SCC papers incorporating different structural model choices.

Figure 2b shows results from the expert survey asking respondents to what extent they agree with the statement that papers incorporating various structural modifications (relative to a baseline model approximating the DICE-2016 IAM (*48*) with a 2020 SCC of around \$40 per tCO_2, similar to the reference distribution in Figure 2a) produce better SCCs than models that omit them. Over 50% of experts agree or strongly agree that a structural model choice constitutes an improvement for all modifications, except the incorporation of aversion to model uncertainty or ambiguity. The strongest agreement is on improvements to Earth system modeling, including the integration of climate system tipping points, and the incorporation of limited substitutability between market and non-market goods in the utility function, with some polarization over the issue of whether using explicit distributional weights, as applied in the literature, improves SCC estimates.

3 Drivers of Variance in SCC Estimates

Figure 1 documents wide variance in published SCC estimates. The rich set of covariates we record allow us to investigate how many different features of SCC modeling - including structural model features, parametric uncertainty, and other model covariates - affect SCC values. Figure 3a shows coefficient values from three regression models using different variation in the dataset. Model 1 compares values across the full SCC distribution (shown in Figure 1). Model 2 adds paper fixed effects, meaning the comparison is between SCC values using alternate modeling and parametric assumptions within the same paper (note both models control quadratically for SCC year, but these coefficients are not shown in Figure 3). Model 3 uses variation between reported central SCC estimates and, where available, comparable "Base SCC" values. Base SCCs were recorded where papers reported values from runs with an unmodified version of the selected IAM. This comparison implicitly controls for SCC year, discount rate, and socio-economic scenario.

Looking first at the effects of structural model choices, Figure 3a reaffirms the importance of persistent/growth damages for SCC estimates represented in the literature across all three regression models. The question of whether and how damages from climate shocks persist in the economy is of major importance for the aggregate costs of climate change, potentially shifting the SCC by an order of magnitude. Consistent and statistically significant increases in the SCC are also seen with the incorporation of explicit distributional weights across space, reflecting the general observation that climate impacts are likely to be regressive in nature (*49, 50*). Allowing for learning over time (typically about equilibrium climate sensitivity or the damage function) tends to decrease the SCC. This is consistent with theoretical models showing that the additional emissions allowed by laxer climate policy can provide a more informative signal about uncertain parameters and lead to better future climate policy (*24*). The effects of other structural model choices are smaller, more mixed, or have larger uncertainties. The estimated effects of changes to the Earth system model or incorporating climate system tipping points tend to be small, a finding consistent with other recent work (*19*). Adding non-linear, stochastic tipping points to damages tends to increase the SCC (Models 2 and 3). Incorporating Epstein-Zin preferences that disentangle risk and time preferences, and allowing for limited substitutability of non-market goods, also tend to increase the SCC, but with uncertainty ranges overlapping zero.

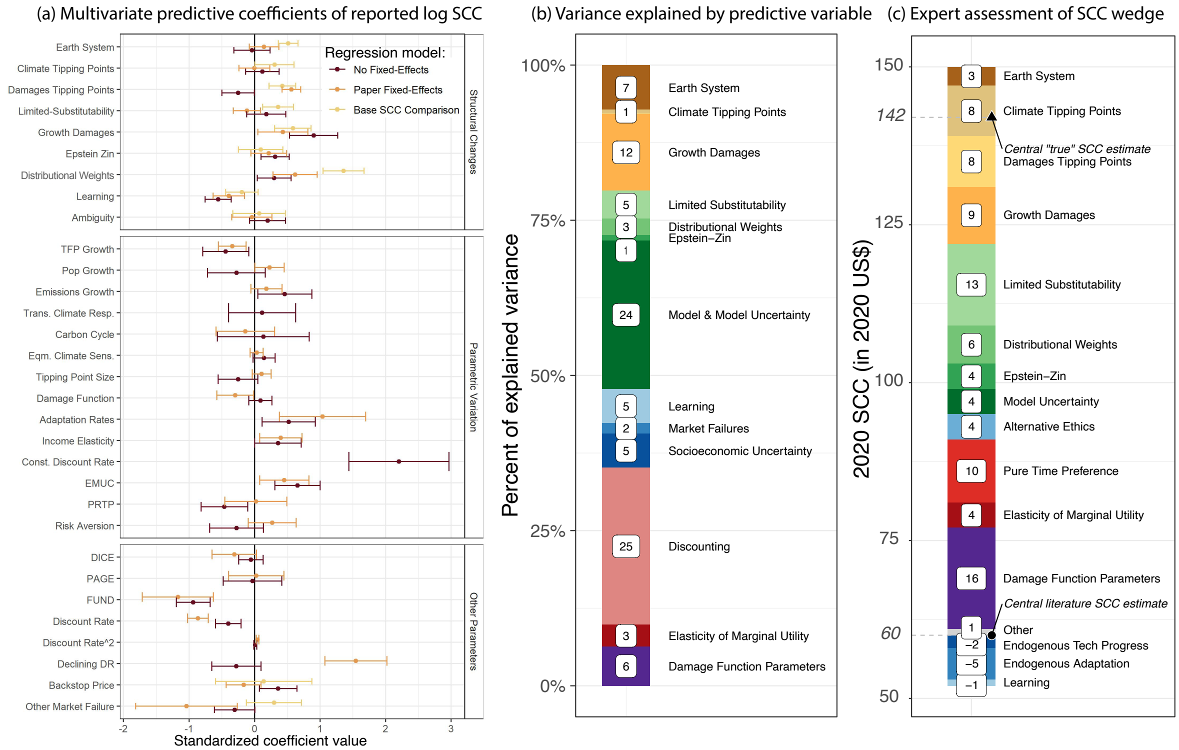

Figure 3: **The relative importance of structural model choices for the SCC.** (a) Regression coefficients for three different multivariate regressions with 95% confidence intervals. The top panel is for our set of 9 structural model choices, the middle panel is for parametric variation, and the bottom panel is for other changes such as the base model. In all cases the dependent variable is logged SCC (in \$ per ton CO_2). Models also include non-linear controls for SCC year (details in Methods). Reported coefficients may be missing if the variable is collinear with the fixed effects or no paper reported a base SCC. Full model results are given in the Appendix. (b) ANOVA decomposition of the variance of logged SCC estimates in the literature, corresponding to the paper fixed-effects regression in column (a). (c) Disentangled components of the wedge between expert's estimates of the literature SCC and the "true" SCC (SCC wedge), aggregated across the 45 experts providing disentangled weights.

Figure 3a also shows the effects of parametric variation on the SCC distribution. Even though parameter values are both increased and decreased, some parameters still have an effect on the average SCC, which could arise either through asymmetry in the input distributions or non-linearities in the SCC calculation. Specifically, parametric variation in TFP growth reduces the SCC, while variation in discount rate parameters and two parameters related to climate damages (i.e., the adaptation rate and the income elasticity of damages) tends to increase it. SCC values calculated using the FUND model tend to be systematically lower.

Second, we perform an ANOVA decomposition analysis of the variance of (logged) SCC estimates in the literature. This analysis combines selected parametric and structural predictors into thematic categories (see SI Section E). Figure 3b shows that the single largest driver of the variance is discounting, followed by model and model uncertainty (i.e., this groups together the identity of the IAM, e.g., DICE, FUND, or PAGE, with the model uncertainty/ambiguity structural model choices), persistent/growth damages, and the Earth system representation (i.e., transient climate response, carbon cycle parametrization, equilibrium climate sensitivity, and Earth system structural model). Note that the overall share of the variance explained by discounting and damage-function parameters (i.e., damage function, adaptation rates, and the income elasticity of damages) is only 31 percent, with most of the remainder relating to structural model choices and model uncertainty.

Figure 3c shows findings from the expert survey asking experts to decompose the difference between their central estimate of the 2020 SCC in the literature and their central "comprehensive" estimate, which we call the "SCC wedge". This measure captures factors experts consider to be both material to SCC estimates and under-represented in the published literature. Figure S20 shows SCC estimates as well as the SCC wedge and its determinants at an individual level. Noteworthy is not only the heterogeneity of SCC estimates across experts, but also the considerable heterogeneity in what explains experts' SCC wedges. The median weights across experts on the determinants of the SCC wedge are zero with the exception of damage function parameters, limited substitutability of non-market goods, persistent/growth damages, and climate tipping points, all of which contribute positively to the SCC wedge. Figure 3c aggregates this information on the expert SCC wedge. We find that damage-function and discounting parameters make up around a third of the SCC wedge. Around two thirds of the SCC wedge is driven by structural model choices, particularly limited substitutability of non-market goods (13%), persistent/growth damages (9%), tipping points in the climate system (8%) and in damages (8%), and distributional weights (6%). Endogenous technological progress and adaptation as well as learning pull down experts' true SCC compared to their estimate of the literature SCC.

Finally, we fit a random forest model to the literature SCC distribution (see SI for details). Machine-learning models can complement linear regressiong and ANOVA by flexibly capturing potential non-linearities and interaction effects in settings with many explanatory variables, optimizing for out-of-sample predictive power. Figure S22 gives an estimate of the importance of different predictor variables based on the change in class impurity associated with that variable (i.e., a variable is of high importance if inclusion in the model substantively increases the discrimination between different SCC values). In addition to the discount rate and SCC pulse year, we see important effects of several structural model choices that also appear in Figure 3, specifically persistent/growth damages and the Earth system model. Also highlighted in Figure S22 is the important role of the damage function. Publication year is the third most important variable, reflecting the general upward trend in SCC values over the 20-year publication period of the sample (*38*).

4 Synthetic SCC Distributions Accounting for Structural Drivers and Parametric Uncertainty

Combining the predictive capabilities of the random forest model, information from the literature on unexplained parametric variation, and expert assessment of model structure and discount rates allows us to develop a 'Synthetic SCC' distribution that integrates over both parametric and structural uncertainty in SCC estimation. We generate distributions of one million SCC values for each emissions pulse year based on random draws that sample from three inputs: (i) the distribution of social discount rate recommendations from a previous expert survey (*51*) (median of 2% and an inter-quartile range of 1-3%); (ii) the distribution of expert assessments of whether different structural model choices in the literature improve the estimation of the SCC, shown in Figure 2b (accounting for covariance between different structural model choices, to the extent that such covariation is reflected in the survey results, using a Bayesian model described in the SI and shown in Figure S21); and (iii) the residual uncertainty not explained by the random forest model. These random draws from expert assessments (i and ii) serve as inputs to the random forest model, to generate predicted SCC values under combinations of model structure and discounting not directly observed in the literature. Residual variation from the random forest model is sampled and added to the predicted values to generate the full synthetic SCC distribution.

Figure 4 and Table 1 report 'Synthetic SCC' distributions as predicted by this model, focusing on an emissions pulse in 2020. Figure 4 starts with the synthetic SCC prediction for a base DICE model and sequentially adds changes before arriving at the full synthetic SCC. Compared with the literature distributions in Figure 1, which could be influenced by factors such as the number of researchers working on certain topics or the ease with which certain types of parametric uncertainty are represented, the Synthetic SCC distribution has a clearer interpretation: it captures epistemic (and normative) uncertainty among experts over model structures and the discount rate, as well as parametric uncertainty represented in the literature and model residuals.

The mean of the 2020 Synthetic SCC distribution is \$467 per ton CO_2, with a median of \$298, substantially higher than either the literature or expert assessments shown in Figure 1. The decomposition shown in Figure 4 shows the contribution of various components to these higher values, showing the substantial contribution of parametric uncertainty (particularly in the damage function and equilibrium climate sensitivity), alternate model structures (particularly alternate earth system models and growth damages), and lower discount rates.

One of the advantages of the random forest model trained on the literature is that it can provide sythetic

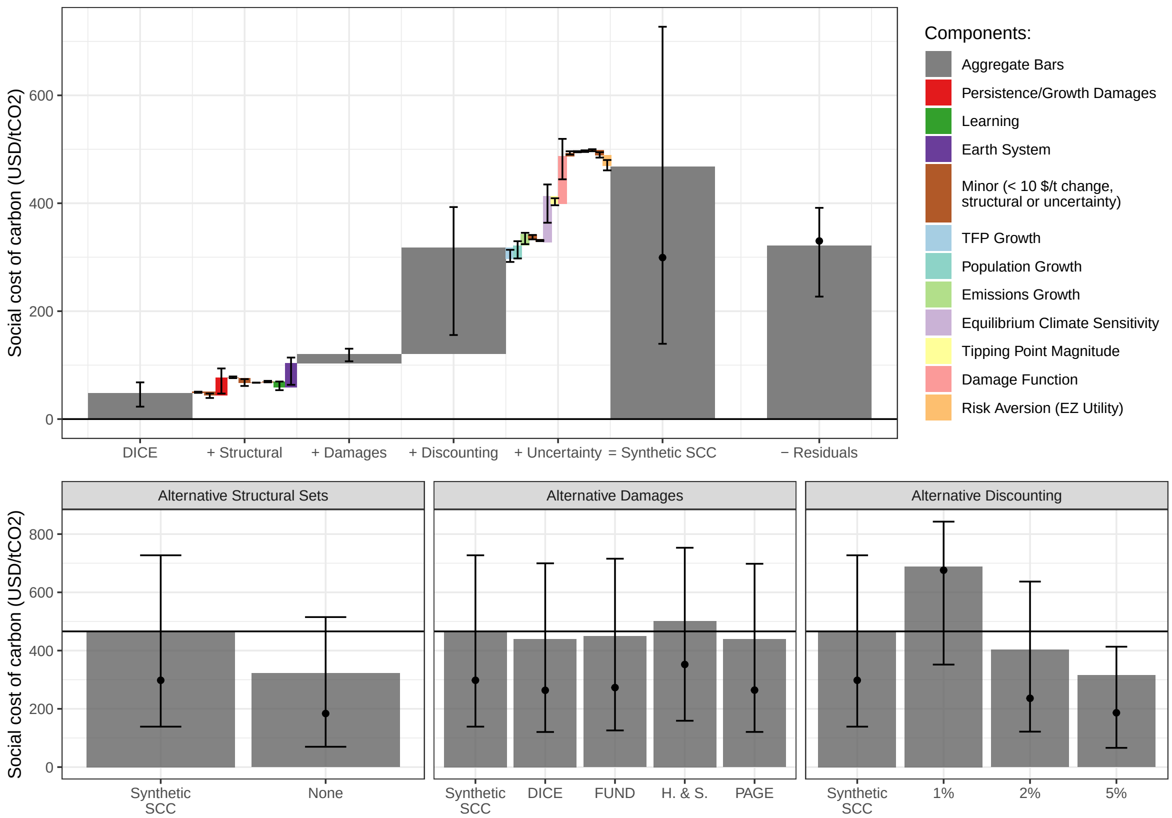

Figure 4: **2020 Synthetic SCC distributions with structural and parametric uncertainty** Top: The 2020 Synthetic SCC distribution arising from the random forest model trained on published SCC values combined with input distributions over discount rates and structural model choices from expert surveys, including its evolution from a baseline DICE comparison. Boxplots show the mean (bar), median (dot), and inter-quartile ranges (error bars). Bottom: Additional distributions of SCCs from the random forest model under different assumptions. From left to right, (a) different structural assumption sets, (b) different damage function assumptions, (c) different constant discount rates, and (d) different pulse years. H&S is the Howard and Sterner (2017) damage function.

SCC estimates under a range of alternate specifications. In the lower portion of Figure 4, we show SCC distributions from the random forest model for different (a) sets of structural model assumptions, (b) damage-functions, (c) social discount rates, and (d) pulse years. Removing all nine structural model choices would reduce the median SCC by more than $100, comparable to increasing the social discount rate to 5%. The Synthetic SCC is relatively insensitive to changes in damage-function type, although this may be because damage-function type is a noisy proxy for damages uncertainty. Reducing the discount rate to 1% would lead to a much higher median SCC of $691.

5 Discussion and Conclusion

We present the most comprehensive synthesis to date of SCC estimates, as well as their parametric and structural drivers. Our study is informed by two complementary data-generating processes: a meta-

	2.5%	5%	10%	25%	50%	75%	90%	95%	97.5%	Mean
	5	14	36	139	298	727	920	1379	1670	467
Structural modifications										
None	-0	3	15	71	185	515	841	884	1131	323
All	-3	2	16	146	325	719	848	1219	1601	452
Discount rates										
1%	24	86	163	356	677	844	1284	1658	2010	692
2%	7	17	36	124	237	638	882	1362	1419	405
5%	1	5	15	65	185	405	847	957	1370	313
Damage functions										
DICE	4	10	27	120	262	700	891	1330	1624	439
FUND	4	10	30	126	275	719	903	1364	1639	452
H. & S.	9	22	51	158	350	755	992	1404	1741	502
PAGE	4	10	29	121	264	704	889	1335	1626	441

Table 1: **Synthetic SCC values.** Distribution of 1 million predictions for Synthetic SCCs (for an emissions pulse in 2020, in 2020 US dollars, rounded to full numbers) from the random forest model, sampling over structural model choices and discount rate ranges from expert surveys. The first row corresponds to the main meta-analytic result. Residual parametric uncertainty is included by randomly sampling from the distribution of random forest residuals. Additional details in Methods.

analysis of published SCC estimates from two decades of research, and an expert survey. Standard discounting and damage-function parameters continue to be important. They explain around a third of the variance in published SCC estimates, and account for around a third of the wedge between what experts estimate is the central value of the SCC in the literature and what they estimate is the "true" SCC, all things considered. But we find that SCC estimates are also strongly shaped by structural model variation in the literature, reflecting alternate characterizations of preferences, Earth system processes, and the nature of climate change impacts. We find particularly important roles for improvements to Earth system modeling, allowing for persistent climate damages, and the representation of limited substitutability of non-market goods. Collectively, these structural model choices explain most of the remaining two thirds of the variation in published SCC estimates and in the experts' SCC wedges.

We find that the distribution of 2010-2030 published SCC values have a median value of $41 per tCO_2 and a mean of $151 (after truncating the upper and lower 0.1% of the distribution). Two further estimates of the overall SCC and its distribution show high probabilities of SCC values well above current official US Government estimates of around $50 per tCO_2. The first comes from experts' estimates of the comprehensive SCC, with a median SCC of $85 and a mean of $160 tCO_2. The second combines the results of a random forest model fitted on the distribution of published SCCs with the results of the expert survey on the SCC, and a separate expert survey on the social discount rate. The 'synthetic SCC' distribution derived from this approach is based on expert assessment of both the discount rate and alternate model structures, but is calibrated to variation in the published literature and also includes residual parametric uncertainty derived from the literature. This distribution has a median value of $298 for the 2020 SCC and a mean of $467 tCO_2 (5%–95% range: $14–$1379). Expert's comprehensive assessment is thus around three times as large as what they deem reflected in the literature, while the Synthetic SCC is yet around three times as large as the mean in the literature, which is strongly driven by uncertainties. Furthermore, the mean estimates from the comprehensive expert assessment and the Synthetic SCC are roughly three and nine times larger than the current formal US SCC estimate, respectively. The Synthetic SCC is also considerably higher in magnitude than other recent estimates that included fewer structural models and a smaller set of parametric uncertainties (*3, 12, 38*), while experts' comprehensive SCC estimate is only slightly lower. There is thus a substantial and varied body of evidence pointing towards using higher SCCs in policy-making. Qualitative evidence from expert recommendations (detailed in SI Section S.2.2.8) and our quantitative findings suggest that future work on the SCC should more closely consider the role of structural model variation alongside parametric uncertainty.

Beyond the multiple lines of evidence for relatively high central estimates of the SCC, our findings also reinforce the critical role of uncertainty in climate policy analysis. All SCC distributions reported in this paper are both broad (with 90% confidence intervals spanning 2 orders of magnitude) and substantially right-tailed (with mean values between 50 and 100% larger than the median). Our analysis of variance suggests this uncertainty arises both from well documented parametric uncertainties (for instance in the equilibrium climate sensitivity) but also from more fundamental structural uncertainties in the nature of the welfare loss from climate change. Median SCC values from the literature, expert assessment and synthetic SCC distribution are all higher - in some cases very substantially so - than the official U.S. government estimate, implying an urgent need to update these estimates (and the corresponding level of climate policy ambition) to reflect current scientific and economic understanding. However, it is also the case that a substantial fraction of the *expected* (i.e. mean) SCC arises from the existence of unlikely but very adverse outcomes. This in turn implies an important role for the insurance-like benefits of climate policy in limiting exposure to downside risk (*8, 35*).

Acknowledgements

We are very grateful to our many survey respondents, to David Anthoff, Kenneth Gillingham, Frikk Nesje, and seminar audiences at AERE 2022, AURÖ 2023, CESifo 2023, University of Potsdam, University of California San Diego, and at PIK for helpful comments, and to Johanna Darmstadt, Luc Esprabens, David Lucius, Nele Steinbrecher, Henry Williams, Angela Zeng, and particularly to Mark Lustig, for excellent research assistance.

Funding

F.C.M. acknowledges support of the National Science Foundation (NSF) award no. 1924378. M.D. gratefully acknowledges support from the DFG under Germany's Excellence Strategy (EXC 2037 and CLICCS) project no. 390683824, contribution to the Center for Earth System Research and Sustainability (CEN) of Universität Hamburg.

Author contributions:

All authors designed the literature review and participated in data collection. FCM and JR lead the data analysis of the literature values and created figures 1, 2 3a-b, and 4. MAD lead development and administration of the expert survey and analysis of results, including Figure 3c. All authors contributed to the writing

Competing interests:

Authors declare that they have no competing interests.

Data and materials availability:

All data, code, and materials from the literature analysis will be made available upon publication. Expert survey data is only available in an anonymized format; this allows producing all main figures and results with the sole exception of supporting analyses that draw on data merged with expert characteristics in SI Sections S.2.2.4 and S.2.2.5.

References

1. N. Stern, *The Economics of Climate Change: The Stern Review* (Cambridge University Press, Cambridge, 2006).
2. J. E. Aldy, M. J. Kotchen, R. N. Stavins, J. H. Stock, *Science* **373**, 850 (2021).
3. K. Rennert, *et al.*, *Nature* **610**, 687 (2022).
4. N. Stern, *Journal of Economic Literature* **51**, 838 (2013).
5. R. S. Pindyck, *Journal of Economic Literature* **51**, 860 (2013).
6. National Academies of Sciences, Engineering, and Medicine and others, *Valuing climate damages: updating estimation of the social cost of carbon dioxide* (National Academies Press, 2017).
7. G. Wagner, *et al.*, Eight priorities for calculating the social cost of carbon (2021).
8. M. L. Weitzman, *Journal of Public Economic Theory* **14**, 221 (2012).
9. A. Otto, B. J. Todd, N. Bowerman, D. J. Frame, M. R. Allen, *Environmental Research Letters* **8**, 024032 (2013).
10. S. Dietz, F. van der Ploeg, A. Rezai, F. Venmans, *Journal of the Association of Environmental and Resource Economists* **8**, 895 (2021).
11. K. Ricke, L. Drouet, K. Caldeira, M. Tavoni, *Nature Climate Change* **8**, 895 (2018).
12. M. C. Hänsel, *et al.*, *Nature Climate Change* **10**, 781 (2020).
13. Y. Cai, T. S. Lontzek, *Journal of Political Economy* **127**, 2684 (2019).
14. B. Crost, C. P. Traeger, *Nature Climate Change* **4**, 631 (2014).
15. K. D. Daniel, R. B. Litterman, G. Wagner, *Proceedings of the National Academy of Sciences* **116**, 20886 (2019).
16. W. Nordhaus, *Proceedings of the National Academy of Sciences* **116**, 12261 (2019).
17. D. Lemoine, C. P. Traeger, *Nature Climate Change* **6**, 514 (2016).
18. Y. Cai, K. L. Judd, T. M. Lenton, T. S. Lontzek, D. Narita, *Proceedings of the National Academy of Sciences* **112**, 4606 (2015).
19. S. Dietz, J. Rising, T. Stoerk, G. Wagner, *Proceedings of the National Academy of Sciences* **118** (2021).
20. I. Rudik, *American Economic Journal: Economic Policy* **12**, 340 (2020).

21. D. Lemoine, C. Traeger, *American Economic Journal: Economic Policy* **6**, 137 (2014).

22. L. Berger, J. Emmerling, M. Tavoni, *Management Science* **63**, 749 (2017).

23. D. Lemoine, C. P. Traeger, *Journal of Economic Behavior & Organization* **132**, 5 (2016).

24. D. Lemoine, I. Rudik, *Annual Review of Resource Economics* **9**, 117 (2017).

25. F. C. Moore, D. B. Diaz, *Nature Climate Change* **5**, 127 (2015).

26. S. Dietz, N. Stern, *The Economic Journal* **125**, 574 (2015).

27. E. J. Moyer, M. D. Woolley, N. J. Matteson, M. J. Glotter, D. A. Weisbach, *The Journal of Legal Studies* **43**, 401 (2014).

28. F. C. Moore, U. Baldos, T. Hertel, D. Diaz, *Nature communications* **8**, 1 (2017).

29. A. Rode, *et al.*, *Nature* **598**, 308 (2021).

30. F. Dennig, M. B. Budolfson, M. Fleurbaey, A. Siebert, R. H. Socolow, *Proceedings of the National Academy of Sciences* **112**, 15827 (2015).

31. D. Anthoff, J. Emmerling, *Journal of the Association of Environmental and Resource Economists* **6**, 243 (2019).

32. A. Rezai, F. Van der Ploeg, *Journal of the Association of Environmental and Resource Economists* **3**, 493 (2016).

33. T. Sterner, U. M. Persson, *Review of Environmental Economics and Policy* (2020).

34. B. A. Bastien-Olvera, F. C. Moore, *Nature Sustainability* **4**, 101 (2021).

35. M. L. Weitzman, *Climate Change Economics* **1**, 57 (2010).

36. M. A. Drupp, M. C. Hänsel, *American Economic Journal: Economic Policy* **13**, 168 (2021).

37. R. S. J. Tol, *Energy Policy* **33**, 2064 (2005).

38. R. S. Tol, *Nature Climate Change* pp. 1–5 (2023).

39. T. Havranek, Z. Irsova, K. Janda, D. Zilberman, *Energy Economics* **51**, 394 (2015).

40. J. Kikstra, *et al.*, *Environ. Res. Lett.* **16** (2021).

41. K. Gillingham, *et al.*, *Journal of the Association of Environmental and Resource Economists* **5**, 791 (2018).

42. S. K. Rose, D. B. Diaz, G. J. Blanford, *Climate Change Economics* **8**, 1750009 (2017).

43. R. S. Pindyck, *Journal of Environmental Economics and Management* **94**, 140 (2019).

44. P. H. Howard, D. Sylvan, *Institute for Policy Integrity* pp. 438–441 (2015).

45. M. J. Schauer, *Environmental and Resource Economics* **5**, 71 (1995).

46. M. A. Drupp, F. Nesje, R. C. Schmidt, *American Economic Journal: Economic Policy* **forthcoming** (2023).

47. D. Anthoff, R. S. J. Tol, Testing the Dismal Theorem (2020).

48. W. Nordhaus, *American economic journal: economic policy* **10**, 333 (2018).

49. S. Hsiang, *et al.*, *Science* **356**, 1362 (2017).

50. M. Burke, S. M. Hsiang, E. Miguel, *Nature* **527**, 235 (2015).

51. M. A. Drupp, M. C. Freeman, B. Groom, F. Nesje, *American Economic Journal: Economic Policy* **10**, 109 (2018).

Supplementary Materials

Moore et al.

September 10, 2023

Contents

S1 Dataset summary

Summary statistics are shown in Table S1, but to provide a taste of the data: the median paper has two authors and reports 6 estimates; it calculates the SCC of an emission in 2020, using a pure rate of time preference (PRTP) of 1% and a elasticity of marginal utility of consumption (EMUC) of 1.45; this results in a central SCC value of $71.25.

The following figures describe other features of the published papers dataset collected as part of the meta-analysis.

Unique values

	N	unique
Papers	1823	147
Estimates	1823	1360
Authors	1823	231
Emissions Scenario	1813	88
Socio-Economic Scenario	1702	63
Damage Function Info.	1142	91

Structural model modifications and assumptions

	N	present
Backstop Price?	1823	19
Other Market Failure?	1823	50
Declining Discounting?	1823	72
Market Only Damages	1823	53
Carbon Cycle	1823	359
Climate Model	1823	382
Climate Tipping Points	1823	50
Damages Tipping Points	1823	168
Persistent Damages	1823	122
Epstein-Zin	1823	77
Model Ambiguity	1823	42
Limitedly-Substitution	1823	60
Inequality Aversion	1823	117
Learning	1823	108
Alternative ethics	1823	16

Uncertainty assumptions

	N	present
Parametric sources of uncertainty		
TFP Growth	1823	120
Population Growth	1823	55
Emissions Growth	1823	70
Transient Climate Response	1823	74
Carbon Cycle	1823	95
Equilibrium Climate Sensitivity	1823	464
Tipping Point Magnitude	1823	130
Damage Function	1823	368
Adaptation Rates	1823	41
Income Elasticity	1823	82
Constant Discount Rate	1823	4
EMUC	1823	61
PRTP	1823	43
Risk Aversion (EZ Utility)	1823	7
Uncertainty information		
Extreme limits	1823	348
Tails ($\geq$ 95%)	1823	261
Central uncertainty (< 95%)	1823	276

Summary values

	N	mean	median	min	max
Authors per paper	1823	2.61	2.00	1.00	9.00
Estimates per paper	1823	12.40	6.00	1.00	249.00
SCC Year	1820	2039.62	2020.00	1995.00	2300.00
Central Value ($ per ton CO2)	1701	252.21	71.25	-23.79	75287.61
Reported Base Model SCC (if applicable)	837	107.00	43.74	-1.98	15063.81
Constant Discount Rate (%)	493	2.92	3.00	0.00	10.00
PRTP	1298	1.01	1.00	-1.10	5.00
EMUC	1234	1.41	1.45	0.00	5.00
RRA	78	6.84	9.50	0.00	10.00
IES	76	1.26	1.50	0.50	2.00
Tail uncertainty level	363	91.19	95.00	50.00	99.90

Table S1: Summary statistics of the SCC dataset. The Unique values table list the number of unique

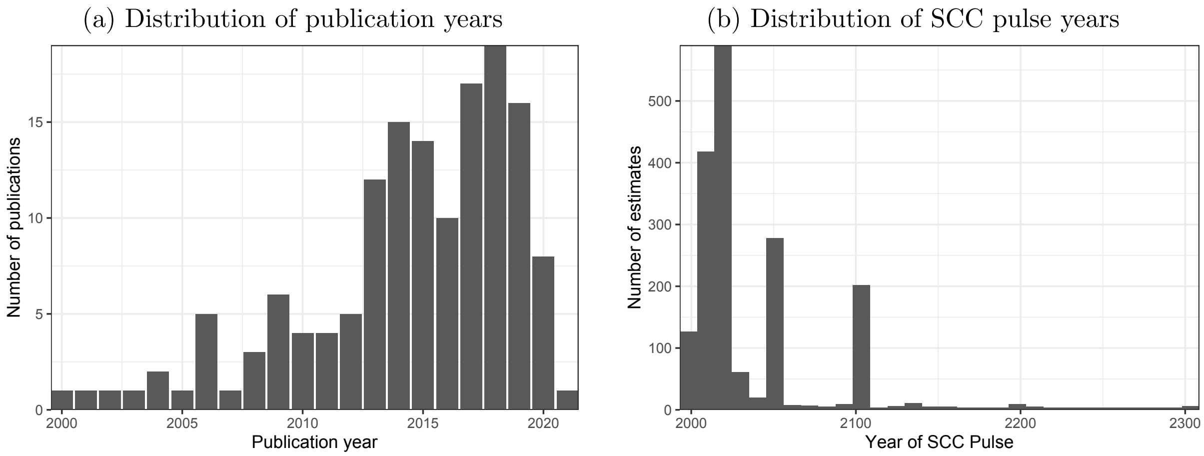

Figure S1: Distribution of (a) publication years and (b) SCC pulse years in the dataset. The distribution of publication years is per publication, while the distribution of SCC pulse years is per estimate since a publication may have multiple estimates.

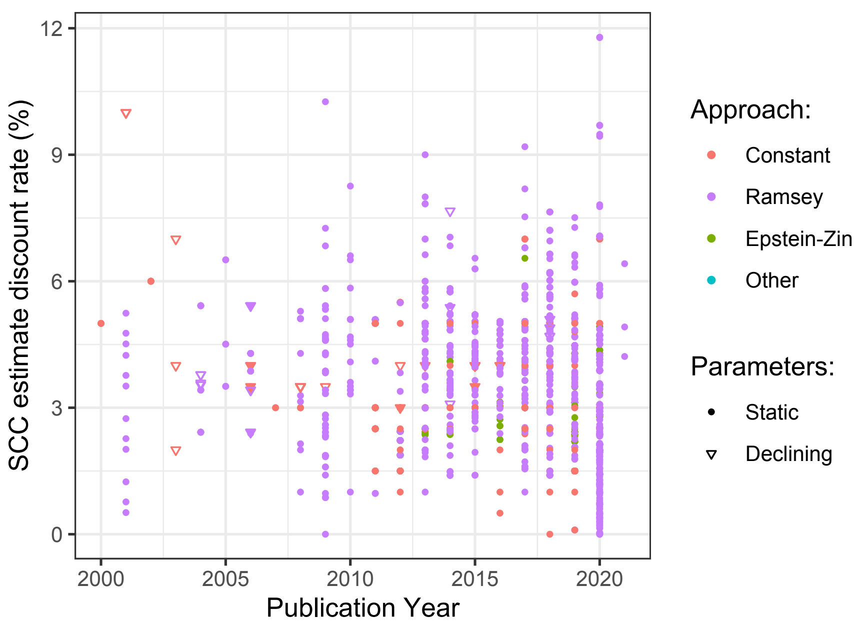

Figure S2: Discount rates used in SCC studies, by year.

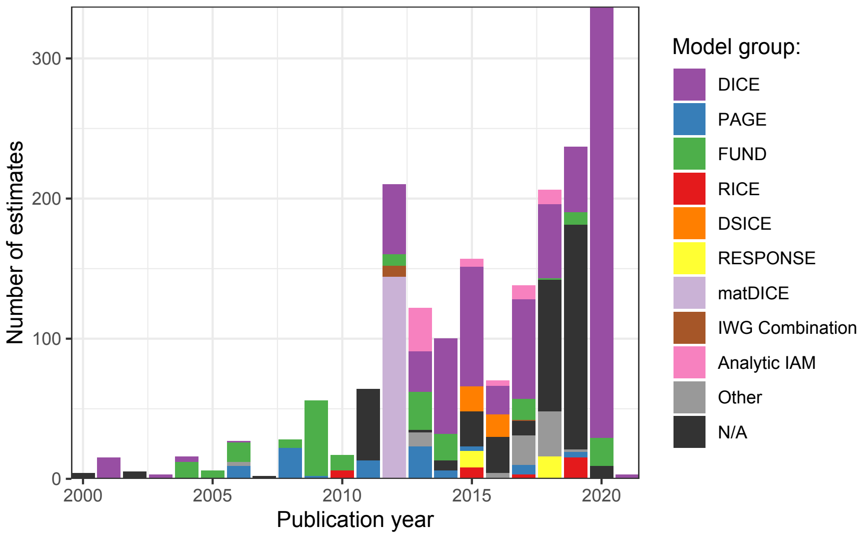

Figure S3: Integrated assessment models used in SCC studies, by year.

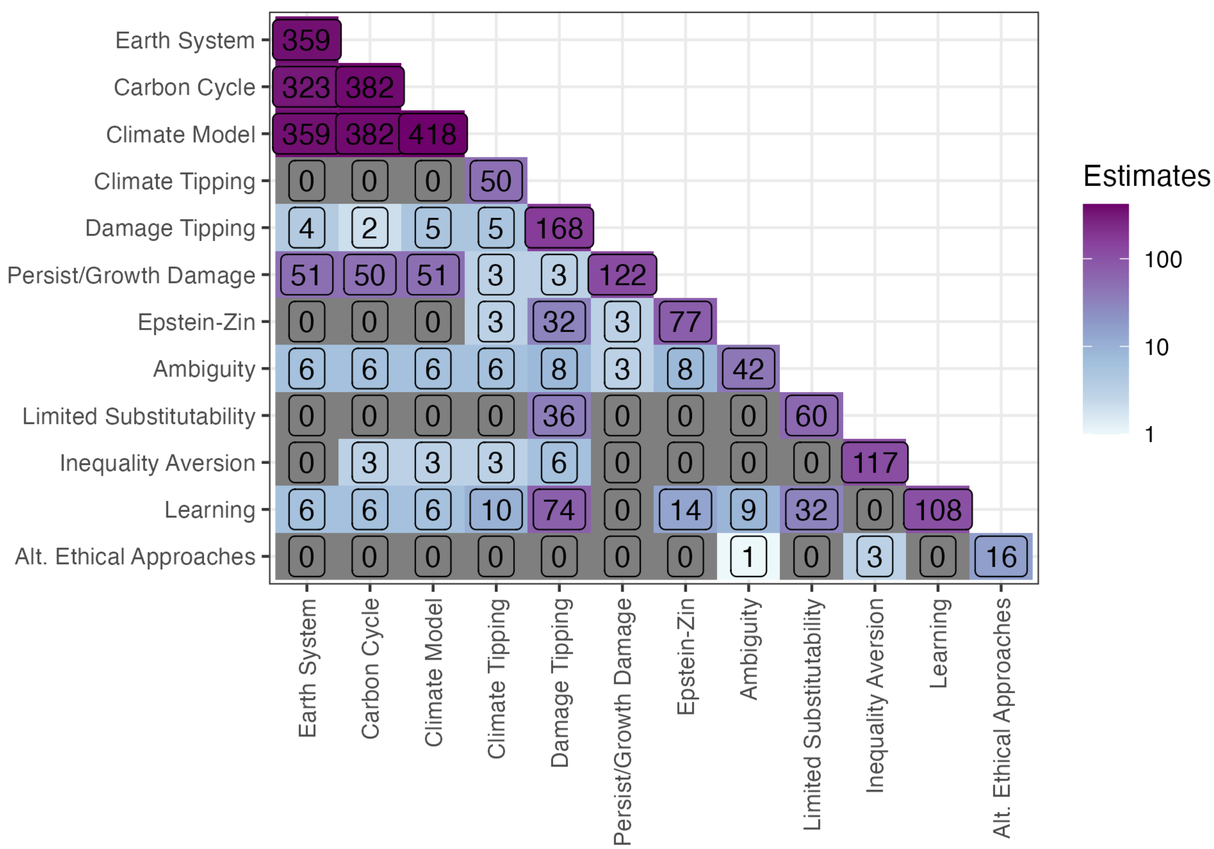

Figure S4: Number of estimates with each pair of structural model modifications. Numbers along the diagonal show the total number of estimates containing each change.

S2 Methods

S.2.1 Meta-analysis

S.2.1.1 Abstract Search

SCC values for use in the meta-analysis were identified from a systematic search of Web of Science, EconLit and Scopus databases. Criteria for the search were peer-reviewed papers published from 2000 to 2020 and containing one of the following search terms:

social cost of carbon, social cost of CO_2, social cost of greenhouse gases, social cost of GHG, optimal climate policy, optimal carbon price, optimal CO_2 price, optimal carbon tax

The search was conducted at the end of September 2020 and so included papers published by that end point. After removing duplicate entries, the search returned 2839 papers. These were further screened by a team of research assistants who read through the abstracts to determine whether the paper was likely to report an original, global social cost of carbon estimate. RAs were instructed to err on the side of keeping papers in the sample if in doubt to avoid dropping relevant papers. 1110 of the 2472 papers initially identified as not containing an original SCC value were re-evaluated by a second RA, an exercise that produced 98% agreement with the original coding. A further 478 abstracts were re-evaluated a third time by a different RA with 99% agreement with the second round of coding.

After the initial abstract review, 295 papers remained that potentially contained original SCC estimates. These were read by members of the author team and SCC values with details of modelling, preference parameters and uncertainty ranges were coded in an initial round of data collection. The author team identified 139 papers producing original SCC estimates of the 295. A further 8 papers meeting the inclusion criteria were identified at this stage and were also included, bringing the total number of papers included in the analysis to 147.

S.2.1.2 Data Collection and Coding

A challenge of attempting to analyze and compare variation in reported SCC across multiple papers is the variety of scenarios, model structures, and parameter values used in different papers. Authors also take different approaches in presenting results and in sampling and investigating uncertainties. A data coding template was developed to extract data on SCC and modeling covariates in a consistent, flexible, and parsimonious way to allow for comparison of values across papers.

The template developed iteratively during the initial round of paper review by the author team. Once all papers had been coded once, 18 papers were coded for a second time by a different person and SCC distributions compared. Using experience from the initial coding and comparing discrepancies and ambiguities arising from the re-coding exercise, we developed a finalized code book describing how SCC values and model covariates should be recorded from papers. All papers were re-read and coded using this finalized code book.

The coding process adopted (given in full as an addition to this Supplement) allows for recording unique SCC values from papers for particular years, discounting assumptions, socio-economic and emissions scenarios, damage functions, and model structure. If the paper reports effects of parametric variation on the SCC, this is also recorded (as distribution quantiles or min and max values) along with the nature of the parametric variation reported in the paper. The final round of coding produced 1823 unique SCC values (or distributions) arising from the 147 papers.

S.2.1.3 Data Cleaning and Standardization

Following the systematic collection of raw data from the papers, we undertook a series of steps to make values comparable across papers. Firstly, SCC dollar values were adjusted to 2020 dollars using the GDP Implicit Price Deflator from the St Louis Fed (*1*) . If a paper did not report the dollar year of the SCC, we first attempted to infer a dollar year based on that used by the baseline model modified or re-calibrated in that paper (e.g DICE2016 or FUND3.9). If this was also unavailable then we assumed a dollar year of 5 years prior to the publication date.

A second important standardization involved imputing comparable discount rates across all values. Approximately 30% of our entries use a constant discount rate to calculate the SCC. The vast majority of the remainder use Ramsey discounting, which depends on two preference parameters (the pure rate of time preference and the elasticity of marginal utility of consumption) and the consumption growth rate. To infer an effective discount rate for SCC values using the Ramsey rule, we merge in information on the consumption growth rate for the relevant time-period under the socio-economic scenario used in the paper. Consumption growth rates from 2020 to 2200 (if available) were identified for multiple different integrated assessment models (10 versions of DICE, RICE 2010, 13 versions of FUND, the SSP scenarios and the older SRES scenarios). After merging in the per-capita consumption growth rate, we calculate the effective discount rate for that SCC value using the Ramsey rule given the reported preference parameters. If we are unable to match a consumption growth rate to a particular SCC value, we impute an estimate based on the average consumption growth rate across all scenarios for that SCC year.

S.2.1.4 Distribution Fitting

We record quantiles of the probability distribution for each SCC value, to the extent that this information is provided by the underlying papers. The full set of quantiles recorded across paper consist of 0.1%, 1%, 2.5%, 5%, 10%, 17%, 25%, 50%, 75%, 83%, 90%, 95%, 97.5%, 99%, 99.9%. Where SCC sensitivity to non-probabilistic parameter changes are reported, we record the minimum and maximum of these. The number of SCC observations reporting each of these quantiles is shown in table S2.

Quantile	Count	Percent
Min	309	17%
0.1th	3	0.2%
1th	4	0.2%
2.5th	4	0.2%
5th	224	12.3%
10th	23	1.3%
17th	72	3.9%
25th	21	1.2%
50th	377	20.7%
75th	21	1.2%
83rd	72	3.9%
90th	29	1.6%
95th	257	14.1%
97.5th	4	0.2%
99th	36	2%
99.9th	7	0.4%
Max	324	17.8%

Table S2: The quantiles recorded across SCC distributions. The number of SCCs with information on each quantile is shown in the Count column, and this as a percent of all SCCs is shown in the Percent column. The Min and Max entries are recorded when sensitivity tests are described with no probabilistic information, while the remaining are used when recording probabilistic analyses (e.g. confidence intervals).

Depending on the available quantiles, we fit different probability distribution functions. The following conditions are applied to each observation:

1. If only the central value is given, the SCC is treated as deterministic.
2. If only the minimum and maximum are given with the central value, a triangular distribution over the provided values is used.
3. If other quantiles are given along with a minimum (maximum), the distribution is bottom-coded (top-coded) to this value.
4. If a 0.1% or 99.9% value is given, the distribution is truncated to these values.

5. If a central value (mean or median) and one other (non-truncating) quantile is given, the distribution is assumed to be Gaussian.
6. If a central value (mean or median) and two other (non-truncating) quantiles are given, the distribution is assumed to be either a Skew normal or a exponentially modified normal, whichever produces a better fit to the quantiles.
7. Otherwise, the distribution is assumed to be a mixture of up to k - 2 Gaussians, where k is the number of fitting values (quantiles and the mean SCC value).
8. If cases 2 - 6 are used, an alternative model consisting of a piecewise uniform distribution with weights from the spans between quantiles is tried as an alternative, and the best-fitting distribution is returned. We also fit a left and right tail extending beyond the most extreme reported quantiles, selecting either a Gaussian, triangular, or exponential distribution based on which best fits the quantiles reported above (right) or below (left) the mean.

In cases where no analytical solution exists to the parameters of the distribution, we evaluate the fit of a potential distribution as

$$\text{RMSE} = \sqrt{(\mu - \hat{\mu})^2 + \sum_k (a_k - \hat{a}_k)^2}$$

where μ is the reported central value, $\hat{\mu}$ is the distribution mean, a_k is the k$^{\text{th}}$ reported quantile, and $\hat{a}_k$ is the corresponding estimated quantile. This is used both to estimate parameters for distributions and to select the preferred distribution according to the rules above. The distributions selected are shown in figure S5 and in S3.

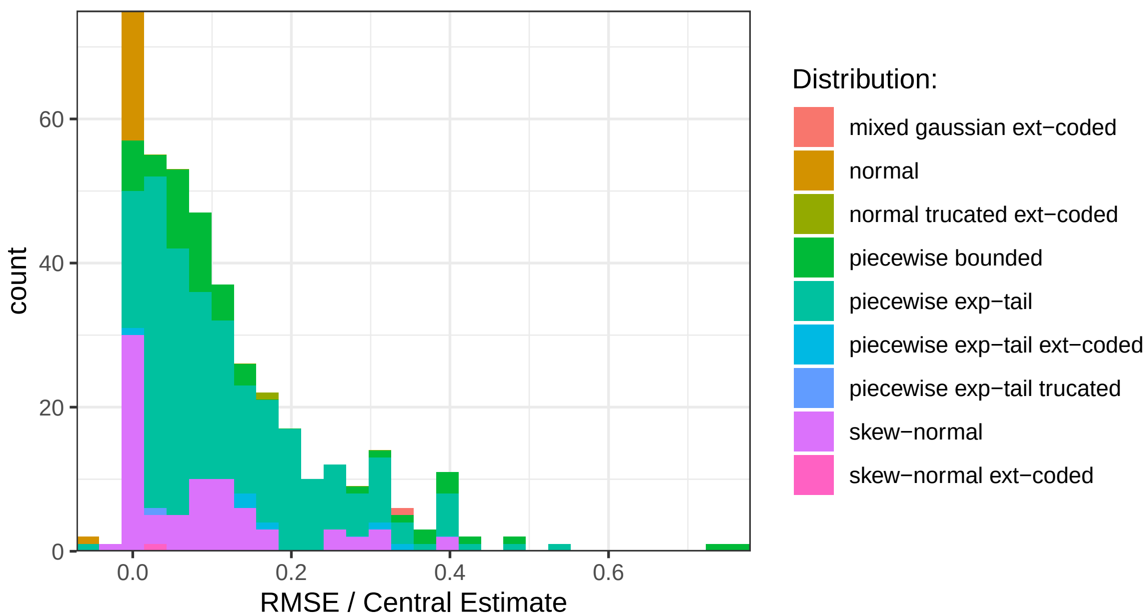

Figure S5: Histogram of the continuous distributions used to fit the quantile information. The solutions are plotted against the degree of miss-fit, described by the ratio of the RMSE to the reported central SCC value.

S.2.1.5 Sampling

Describing the distribution of SCC values in the literature requires sampling over papers and reported values. We investigate three different approaches to sample over the set of SCC papers:

1. **Equal Paper Weighting:** Each paper in the data-set receives equal weighting

(a) Count of distributions fit to data

Distribution	Unbounded	Truncated
Delta	1140	
Triangle	273	
Gaussian	19	1
Skew-Normal	102	2
Piece-wise Uniform	225	61
Total stochastic	683	

(b) Count of tails fit to piece-wise distributions

Distribution	Left-Tail	Right-Tail
Triangle	139	99
Gaussian	88	97
Exponential	3	34
Total piece-wise tails	230	230

Table S3: Information about the distributions fit to the reported SCC information. (a) The general distribution applied, and whether it is truncated or not. Delta, Triangle, and Gaussian distributions are applied in particular cases, while Skew-Normal and Piece-wise Uniform distributions are chosen based on goodness-of-fit. The Total stochastic row excludes Delta distributions. (b) For piece-wise uniform distributions, the tails fit based on quantile information. In total 230 piece-wise uniform distributions have tail information, including 5 which are truncated; the remaining 56 are bounded (uniform to the minimum and maximum).

2. **Informational Weighting:** Papers more likely to contain more independent estimates are weighted more heavily than papers likely to have estimates highly correlated with other papers in the dataset. We operationalize this by calculating a shared co-author index that compares the average number of co-authors shared between paper i and paper j compared to an estimated null value based on 250 random draws that fix the number of authors in the sample and the number of authors on papers, but randomly reshuffles co-authors. Papers with average shared authorship less than the mean of the null distribution receive full weight while those with higher values receive lower weights that gradually decline with higher levels of shared authorship across the 147 papers in the dataset.
3. **Citation Weighting:** Papers with higher citation counts (based on Google Scholar) are weighted more heavily. To avoid mechanically placing higher weight on older papers, weighting is based on the average citations per year since publication

After sampling the paper, using the three alternate sampling schemes, all SCC observations (either single values or distributions, depending on whether the paper reported parametric uncertainty) are equally likely to be sampled. Table **??** gives quantiles from the three alternate sampling schemes and shows relatively small differences across the distributions at the median and lower half of the distribution. But the co-author weighted and citation-weighted sampling schemes have substantially more probability mass in the upper tails, particularly the citation-weighted distribution. For simplicity, analysis and discussion in the paper focuses on the distribution using the equal-weighting of papers.

S.2.1.6 Sensitivity Analysis

Figure S6 shows the change in the mean SCC value for the full literature distribution after iteratively dropping each paper in the dataset (Figure S6 shows the 15 papers with the largest effect on the mean SCC). By far the highest-leverage paper is (*2*), largely because this paper includes a value of over \$65,000 per ton CO_2 under a "super-low" discounting scheme with a constant 0.1% discount rate (Table J-1). Dropping this paper reduces the mean 2010-2030 SCC by \$136. Only two other papers (*3, 4*) change the mean SCC by more than \$10 and one (*5*) changes it by between \$5 and \$10. The remaining 143 papers affect the mean SCC value by less than \$5 each. Because of the extreme effect of this single supplementary value, we drop this most extreme value from (*2*) in all subsequent analysis, including all results reported in the main text.

We further investigate the degree to which the shape of the literature SCC distribution shown in 1 is

Quantile	Equal-Weighting	Co-Author Weighted	Citation Weighted
0.01	-9.6	-10.1	-19.1
0.025	0	1.0	0.2
0.05	2.3	4.1	3.3
0.1	6.5	9.1	7.8
0.25	17.5	22.3	17.8
0.5	46.4	56.4	52.1
0.75	130.0	157.2	177.1
0.9	290.2	354.8	591.3
0.95	615.3	660.6	1046.2
0.975	934.8	1046.2	1495.7
0.99	1438.60	1640.1	3038.44

Table S4: Quantiles of the full SCC distribution using alternate weighting schemes to sample papers. Results in the main text sample uniformly from all papers (Equal-Weighting). Here we compare quantiles from this distribution with alternate weighting schemes. Co-Author Weighted down-weights papers that share large numbers of co-authors with other papers in the dataset on the basis that the information content of these papers may be not fully independent. Citation Weighted weights the sampling of papers based on citation counts, normalized by the number of years since publication. The central 5% - 75% of the distributions are very similar, but citation weighting in particular results in more probability mass on SCC values above \$200 per ton CO_2.

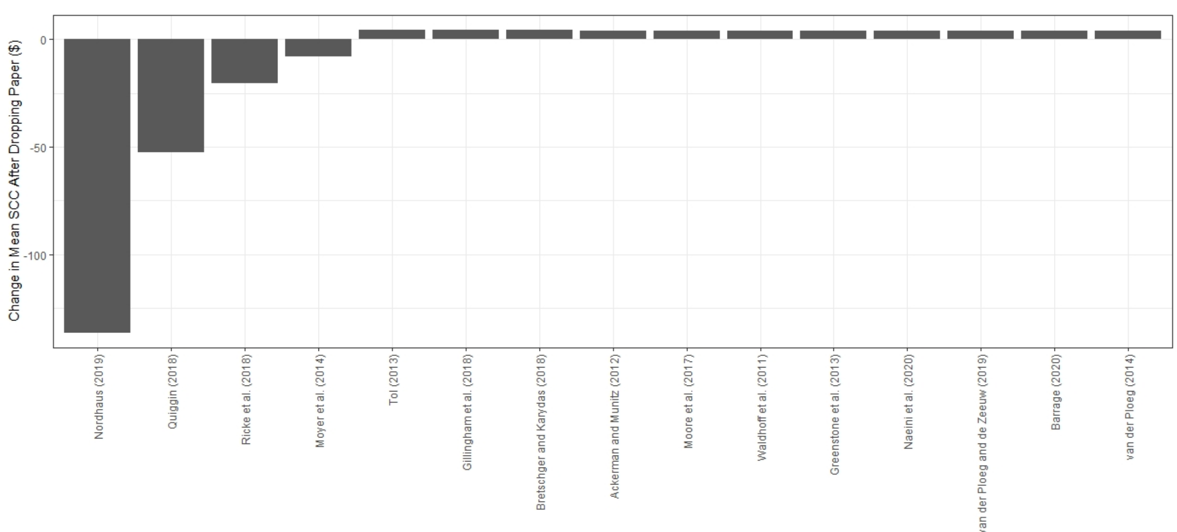

Figure S6: 15 Papers with Largest Effect on Mean SCC for the 2010-2030 period, Defined as the Change in Mean SCC Value after Dropping the Paper from the Distribution

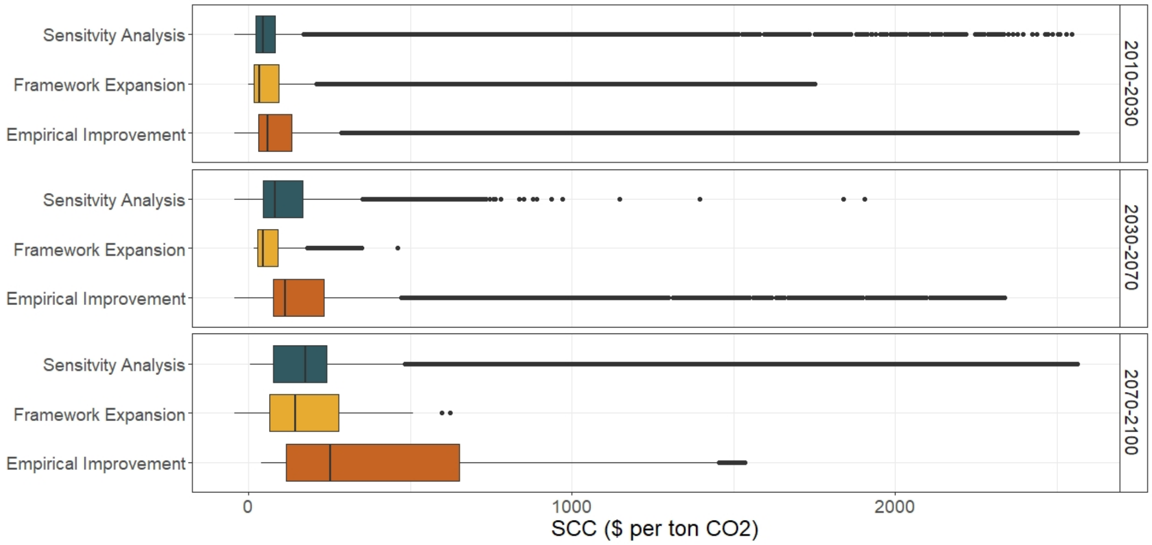

Figure S7: Same as Figure 1 but separated by paper type

driven by papers reporting pure sensitivity analysis rather than papers reporting economic enhancements or empirical improvements to SCC estimates. We split the distribution in Figure 1 into three classes: pure sensitivity analyses, papers expanding the economic framework used to calculate the SCC, and papers reporting empirical improvements to elements of the IAMs supporting SCC calculations. (This information was recorded during original data collection and more specific descriptions of these classes is given in the attached coding document).

Figure S7 shows that high SCC values and the long right tail of values are not driven by papers conducting pure sensitivity analysis, particularly for emissions pulse years in the 2010-2030 period. Instead, the highest mean and median values appear in papers categorized as empirical improvements.

S.2.1.7 Tail Behavior

One of the most notable features of Figure 1 is the long right tail, extending well above $500 per ton of CO_2. The question of the role that low probability but very bad outcomes (i.e., the "right tail" of climate damages) should play in driving climate policy has been written about extensively. In developing his "dismal theorem", Weitzman described the potentially high sensitivity of expected climate damages to behavior in the far tail of the distribution (*6, 7*). In extreme cases, the presence of fat tails may lead to unlimited downside exposure and a distribution with an infinite mean. Even in less extreme cases, substantial probability mass in the tails may cause the expected value to be highly sensitive to necessarily subjective judgements regarding the probability of very bad outcomes (*7*).

Previous work in the literature has shown evidence for a long-right tail in climate damages. Anthoff and Tol (*8*) for example, looking only at parametric uncertainty included in the FUND model, find evidence for fat tails as the mean of the distribution continues to increase with the number of Monte Carlo runs. Other work looking at parametric uncertainty across multiple variables and several IAMs also typically shows a right tail on the SCC distribution (*9, 10*). Recently, Anthoff and Tol (*11*) examined evidence for fat tails in both the parametric distribution of DICE, PAGE and FUND, as well as across published estimates of the SCC, again finding support for fat tails in the FUND model but only mixed evidence in the other models and in the meta-analysis.

The presence of fat tails in a distribution can be tested using the mean excess function (MEF) – the mean of the distribution conditional on being above some threshold (*12–14*). If the mean above the threshold is increasing faster than the threshold itself – then it indicates the presence of a fat tail. The slope of the mean excess function can be used to find the value of the shape parameter for the Generalized Pareto Distribution (GPD) that best fits the distribution (*12*), and the tail index (α) of the distribution. A tail index greater than 2 indicates a thin-tailed distribution, $1 < \alpha < 2$ is a thick tailed distribution with finite mean but infinite variance, while $\alpha < 1$ is a fat tailed distribution with infinite mean and variance. Table S5 shows estimates of the tail index using multiple threshold quantiles for the estimation. Regardless of how we weight observations or set a minimum threshold for inclusion in the sample, we consistently find mean excess function slopes greater than 1, and α values between 1 and 2. This is indicative of a fat-tailed distribution with infinite variance but finite mean.

Table S5: Estimates of the Mean Excess Function slope, the Generalized Pareto Distribution shape parameter, and tail index.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
MEF Slope	1.731***	1.661***	1.315***	1.477***	2.856***	2.746***	2.624***	1.477***
	(0.085)	(0.043)	(0.022)	(0.006)	(0.094)	(0.060)	(0.035)	(0.006)
Num.Obs.	67	135	338	1316	67	135	338	1316
Minimum Threshold Percentile	95	90	75	0	95	90	75	0
Observational Weights	Num. Obs.	Num. Obs.	Num. Obs.	Num. Obs.	1/Num. Obs.	1/Num. Obs.	1/Num. Obs.	1/Num. Obs.
GPD Shape Parameter	0.63	0.62	0.57	0.6	0.74	0.73	0.72	0.6
Estimated Tail Index	1.58	1.6	1.76	1.68	1.35	1.36	1.38	1.68

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note:

Standard errors are robust to heteroskedasticity. All estimates are from a sample that excludes Nordhaus (2019). Columns 1-4 weight observations of the mean excess by the number of SCC observations used to compute it. Columns 5-8 weight observations with the inverse.

S.2.1.8 Damage function based SCCs

For 46% of the SCC observations, we have constructed a simple functional representation of the underlying damage relationships. These consist of 13 versions of DICE damage functions across 241 estimates, 7 versions of FUND across 83 estimates, 2 versions of the Howard & Sterner function across 261 estimates, 5 estimates using PAGE damages, 68 estimates using Weitzman damages, 11 estimates using Dietz & Stern damages, and 168 estimates where another explicit functional form was used. A heat map of these various damages is shown in figure S8.

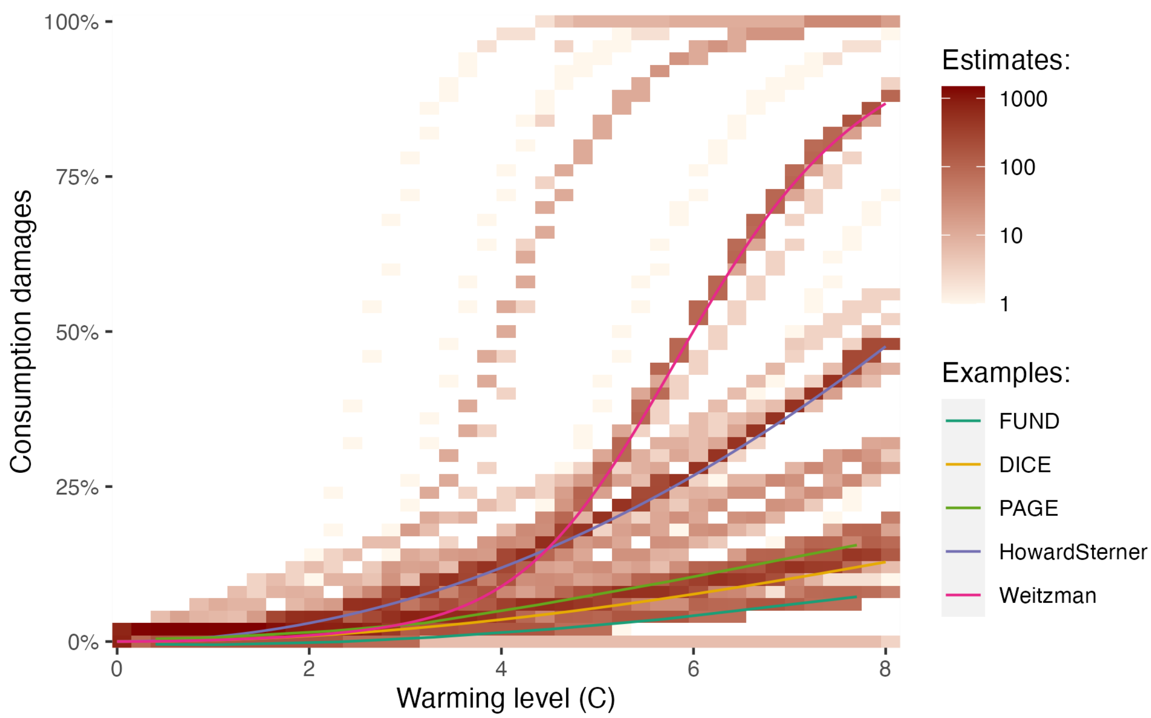

Figure S8: Damages projected across temperatures, shown as a heat map of the occurance of the given damages across studies. Illustrative common damage functions shown as curves.

The SCCs that result from these damage functions are dependent upon not only the structural features of interest in this paper, but also the climate and socioeconomic scenario, discounting, and baseline assumptions.

As a simple proxy for damages, we generate a damage function based SCC under identical conditions. This is done by simulating temperatures under DICE 2013 for a baseline RCP 8.5 scenario and an additional pulse in 2020. Damages are calculated as fractional losses of GDP in each year and totaled under a 3% discount rate. These total damages are translated into dollars per ton assuming a constant global GDP of $84.54 trillion.

Figure S9 shows the damage function based SCC compared to reported SCCs. Reported SCCs show a spread around each damage function based SCC, based on the range of assumptions used beyond the damage function. However, damage function based SCCs raise proportionally with reported SCCs, on average. A regression explaining reported central SCC as a function of SCC year, discount rate, and damage function based SCC achieves an R^2 of 0.52.

Table S6: Odds ratio of being in the top 10 percent of SCC values.

	Odds Ratios for Being in Right Tail of Distribution		
	2010-2030	2030-2070	2070-2100
Parametric Uncertainty			
Adaptation Rates	3.0	0.0	Inf
Carbon Cycle	0.3	0.0	12.1
Constant Discount Rate	0.0	NA	NA
Damage Function	2.9	1.7	5.9
Elasticity of Marginal Utility	0.0	0.0	NA
Emissions Growth	1.3	0.0	NA
Epstein-Zin Risk Aversion	0.0	NA	NA
Equilibrium Climate Sensitivity	1.7	0.7	2.6
Income Elasticity	3.0	0.0	0.0
Population Growth	0.0	0.0	0.0
Pure Rate of Time Preference	1.4	0.0	0.0
TFP Growth	0.5	0.0	1.5
Tipping Point Magnitude	0.3	0.0	0.0
Transient Climate Response	0.0	0.0	NA
Structural Change			
Alternative Ethical Approaches Not Discount	0.0	NA	NA
Ambiguity Model Uncertainty	1.6	3.1	0.0
Climate Tipping	2.1	0.0	0.0
Damage Tipping	0.7	0.0	0.0
Earth System	2.3	0.9	1.4
Epstein-Zin Preferences	0.8	0.0	0.0
Inequality Aversion	1.1	0.0	0.0
Learning	0.2	0.0	0.0
Limitedly Substitutable Goods	0.3	0.0	0.0
Persistent Growth Damages	8.2	3.8	0.0

Note:
The values are the odds ratio of being in the top 10 percent of the SCC distribution in each of the three periods comparing SCCs computed using the listed change to those computed without. NA values correspond to when we do not have SCC values with the listed change during that time period (e.g. emissions growth uncertainty in the late period). 0 values correspond to the SCC never being in the top 10 percent with the listed change.

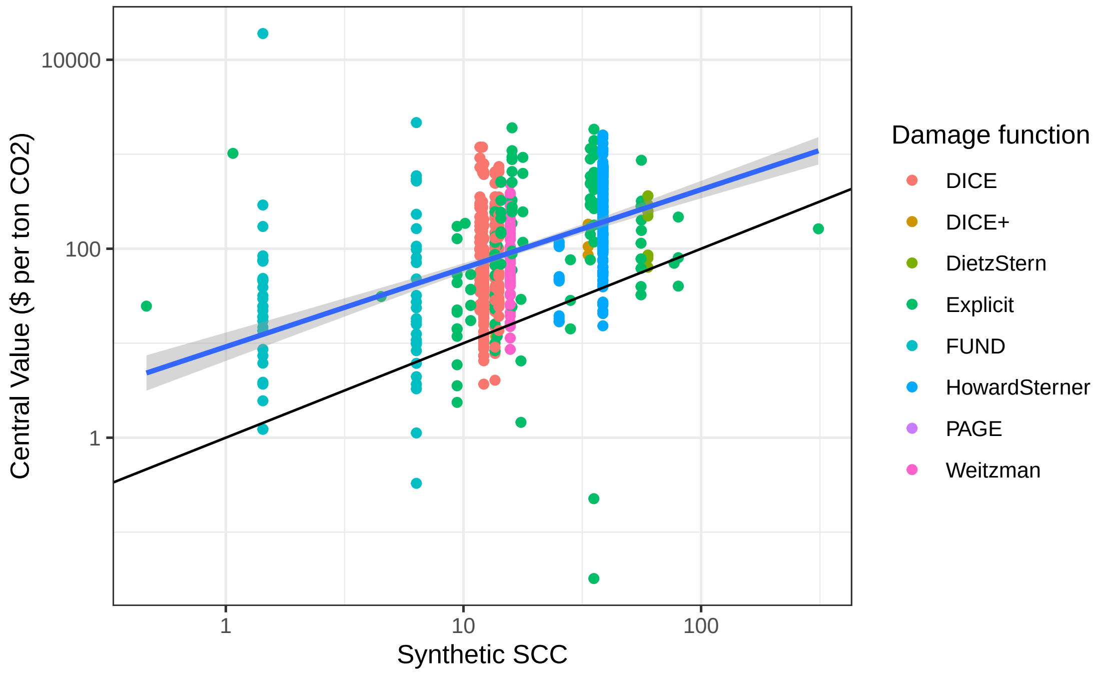

Figure S9: Calculated damage function based SCC values, compared to reported central values, for each estimate for which we can calculate a damage function based SCC. Dots are colored by the class of damage function used, where 'Explicit' refers to an explicitly described functional form rather than a standard damage function. The black line shows has unit slope, and a regression line is shown in blue.

S.2.1.9 Multi-Variate Regression

The multi-variate analysis shown in Figure 3a uses three different types of variation in the data.

1. **Full Variation (No Fixed Effects)** The first model identifies the effects of different modeling decisions using the full variation across the whole SCC distribution. The estimating equation is (betas omitted for clarity):

   $\log(SCC_{dp}) = Year_i + Year_i^2 + DiscountRate_d + DiscountRate_d^2 + \mathbf{Struc_d} + \mathbf{Param_d} + \mathbf{Other_d} + \epsilon_{idp}$

   Where the dependent variable is the log of an SCC observtion from distribution d, from paper p. (Using logs requires dropping the 2% of the distribution below $0.) $\mathbf{Struc_d}$ is set of nine indicator variables indicating whether or not the estimate comes from a model that includes a particular structural model modifications. $\mathbf{Param_d}$ is a set of indicator variables describing whether the estimate is drawn from a distribution containing parametric variation in one of 14 possible parameters. $\mathbf{Other_d}$ contains an additional six binary variables describing the estimate, such as whether it is a backstop price, and whether it is derived from versions of the DICE, PAGE or FUND models. Finally, the specification includes quadratic controls for the SCC year and discount rate. Residuals are clustered at the distribution level (i.e. allowing for correlation in the error term for draws from the same distribution of parametric uncertainty).

2. **Within-Paper Variation (Paper Fixed Effects)** The second approach adds fixed-effects by paper. The specification is the same as above, except for the addition of paper fixed-effects that control for all average differences between papers. Model parameters are then estimated off of variation reported *within* a single paper. For instance, a single paper might report the SCC under alternate discount rates, or with and without various structural model modifications to the model.

3. **Base SCC Comparison** The final multi-variate comparison uses variation between central SCC

values and a standard model comparison point, which we term a "base SCC". Many papers run a standard version of an IAM, make some modification and report the effect of that modification on the SCC. Since the base SCC values are not original estimates, we record them specifically as comparison points, not separate observations of the SCC. The 'Base SCC Comparison' regression specifically uses variation between the base SCC values and other central SCC values that are comparable except for specific structural model modifications reported in the paper.

The equation for this estimate is (betas omitted for clarity):

$$log(SCC_{trs}) = \textbf{Struc} + \textbf{Other} + \theta_{trs} + \epsilon_{trs}$$

Where the dependent variable is the log of the SCC for a ton emitted in year t, calculated using discount rate r, using scenario s, using both original and base SCC values. **Struc** and **Other** are defined as above. θ_{trs} is a fixed-effect for each unique combination of SCC year, discount rate, and scenarios. This means that parameters are only estimated off of variation coming from differences in model structure (or a limited set of other variables that might affect the SCC), conditioning on these other SCC determinants. Note that because we do not have data on parametric variation in the base SCC values from most papers, this regression uses only reported central values, not the full distribution including parametric uncertainty.

Note that the different variation used in these three approaches is distinct and, correspondingly, that the interpretation of the coefficients shown in Figure 3a is different for each model. For example, the effects of incorporating tipping points in the damage function are positive for the Base SCC comparison and the Paper Fixed Effects models, but are slightly negative in the Full Variation model. This can be rationalized if the set of papers that allow for tipping points in the damage function tend to, on average, incorporate them into a base model that produces an SCC on the low end of the full set in the data. This means that, on average, SCC from models with this structural model modification are lower than average (i.e. a negative coefficient in the Full Variation model), but the effect of this change, relative to a model that is the same except for this change, is an increase in the SCC (i.e. positive coefficients in the Paper Fixed Effects and Base SCC Comparison models).

When all three sets of variables are included, a complex regression tree emerges (see Figure **??** a). The inclusion of a Carbon Cycle is the most predictive division, in line with the previous analyses. Beyond this, the SCC year and effective discount rate are highly predictive. Since these variables enter into the SCC continuously, while the regression tree is best at identifying discontinuous changes, we next factor our the analysis variables, as quadratics. These variables explain 18.5% of the variance in log-SCC values.

The regression tree for the remaining variation continues to the dominated by the carbon cycle decision. At the second level, the use of persistent damages is most significant, with models using persistent damages showing double the SCC of the average model when carbon cycle changes are not used, and showing 10 times the SCC when carbon cycle changes are included. Below this level, changes in the damage function reflected in the damage function based SCC are more predictive than any other structural model modification.

S.2.1.10 Analysis of variance

The analysis of variance determines the portion of the total variance attributable to each predictor. For this analysis, we use Paper Fixed-Effects regression discussed in section 3, except that we drop the SCC year quadratic and subset the data to just SCC years from 2010 to 2030 (inclusive). The estimates for each predictor are reported in table S7.

When reporting variances in the bar chart in the main text (figure 3b), parametric and structural predictors are combined. Specifically, Trans. Climate Resp., Carbon Cycle (Param), Eqm. Climate Sens., and Earth System are reported as "Earth System"; Damages Tipping Points and Tipping Point Size are reported as "Damages Tipping Points" (but not labeled in the figure); Epstein Zin and Risk Aversion are reported as "Epstein-Zin"; Model group (an indicator of DICE, FUND, PAGE, or other) and Ambiguity are reported as "Model & Model Uncertainty"; TFP Growth, Pop Growth, and Emissions Growth are reported as "Socioeconomic Uncertainty"; Const. Discount Rate and a quadratic

Source of variation	Estimate	LOO Range		Bag Range	
Earth System	5.05	5.24	[5 - 7]	7.75	[0 - 41]
Climate Tipping Points	0.56	0.58	[0 - 1]	1.24	[0 - 4]
Tipping Point Size	0.06	0.07	[0 - 0]	1.42	[0 - 3]
Growth Damages	12.29	12.9	[12 - 17]	23.45	[5 - 80]
Epstein Zin	0.71	0.77	[1 - 1]	3.39	[0 - 24]
Ambiguity	0.01	0.02	[0 - 0]	0.53	[0 - 1]
Limited-Substitutability	4.53	4.66	[4 - 6]	4.95	[0 - 31]
Inequality Aversion	2.68	2.73	[2 - 4]	2.23	[0 - 9]
Learning	5.49	5.75	[5 - 7]	14.07	[6 - 84]
TFP Growth	2.57	2.74	[2 - 4]	9.76	[1 - 53]
Pop Growth	0.16	0.24	[0 - 1]	6.58	[0 - 36]
Emissions Growth	2.76	2.82	[2 - 4]	1.55	[0 - 5]
Trans. Climate Resp.	0.28	0.3	[0 - 0]	1.18	[0 - 4]
Carbon Cycle (Param)	0.2	0.23	[0 - 1]	0.55	[0 - 3]
Eqm. Climate Sens.	1.7	1.73	[1 - 3]	1.17	[0 - 4]
Damages Tipping Points	0.07	0.07	[0 - 0]	1.68	[0 - 9]
Damage Function	4.52	4.88	[4 - 7]	12.07	[1 - 49]
Adaptation Rates	0.91	1.03	[1 - 2]	7.61	[1 - 43]
Income Elasticity	0.42	0.47	[0 - 1]	4.59	[0 - 15]
Const. Discount Rate	0.59	0.75	[0 - 2]	7.49	[0 - 45]
EMUC	3.47	3.67	[3 - 5]	9.64	[2 - 45]
PRTP	0	0.06	[0 - 1]	5.26	[0 - 46]
Risk Aversion	0.16	0.15	[0 - 0]	0.99	[0 - 6]
Model group	23.59	23.87	[21 - 27]	27	[13 - 69]
Other Market Failure	1.45	1.52	[1 - 2]	1.72	[0 - 7]
Damage-based SCC	0.44	0.56	[0 - 2]	2.82	[0 - 16]
Discount Rate	25.32	25.77	[24 - 29]	37.78	[23 - 80]

Table S7: Analysis of variance (ANOVA) results corresponding to the regression in figure 3b. All entries are percentages of explained variance, excluding variance from paper fixed effects. The Estimate column reports the percent of explained variance across all included variables, as shown in the figure. The LOO Range reports the central estimate and range of variances calculated when each individual predictor is dropped, and the Bag Range reports central estimate and range of variances calculated when a random subset of these predictors is selected. Ranges are 95% ranges.

in (effective) Discount Rate are reported as “Discounting”; and the Damage-based SCC, Damage Function, Adaptation Rates, and Income Elasticity are reported as “Damage Function Parameters”.

S.2.2 Expert survey

We invited all authors of SCC estimates included in our meta-analysis, and for whom we could obtain a workable e-mail address, to participate in the expert survey. In the e-mail introduction (see text in annex to this supplement), we explained that the goal of the survey is to elicit experts' views on the SCC, uncertainty about its value, and how various structural model modifications affect the SCC. We communicated that results will be published without identifying any individual participant, and that we had obtained approval from the research ethics review board at UC Davis and approval from the social science research deanery and social science research laboratory at the University of Hamburg. We sent invitations to the effective population of 176 SCC authors on May 22, 2022, and closed the survey on July 07, 2022. 72 of the invited SCC authors participated, of which 68 provided quantitative responses and 48 responded non-anonymously, with response rates of 41%, 39% and 28%, respectively, which compares very well with similar expert surveys (*15–18*).

S.2.2.1 Survey design

The survey contains four questions with sub-questions and was conducted online via platform SoSci Survey (screenshots with all survey details are provided in figures S10 to S18). The first question elicits estimates of the distribution of the SCC in the literature for the year 2020 (the central value and the 2.5 and 9.75 percentiles), akin to a prediction study as in the experimental literature (*19*), and the distribution of the appropriate or "true" SCC in the year 2020 "all things considered". We consider this as an expert's 'comprehensive estimate'. In the second question, we elicit experts' views on how strongly selected structural model modifications in the literature affect SCC estimates and whether and to what extent these model extension as reflected in the literature provide improved SCC estimates. In the third question, we elicit experts views on what drives the potential wedge between the SCC in the literature and the "true" SCC all things considered. Finally, in the fourth question, we asked for expert's views on the most important steps for improving estimates of the SCC going forward with an open ended qualitative response option.

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CLUSTER OF EXCELLENCE
CLIMATE, CLIMATIC CHANGE, AND SOCIETY (CLICCS)

UC DAVIS
UNIVERSITY OF CALIFORNIA

Introduction and Purpose

This is an expert survey on the social cost of carbon (SCC); i.e., the marginal damage cost of emitting a metric ton of CO2. You have been invited to participate because you have authored a peer-reviewed study estimating the SCC published between 2000 and 2020. This short survey contains four questions and may be completed in around 15 to 20 minutes.

Our survey aims to elicit authors' views on the SCC and 'structural' changes to Integrated Assessment Models (IAMs) that affect the SCC. Such structural changes – including to the earth system module(s), climate damage function(s), and utility/welfare functions – may increase or decrease the SCC and are explored selectively in the literature, rather than systematically.

We expect experts' views on the importance of these structural changes to vary, just as we expect experts to differ in their assessment of the importance of other determinants of the SCC, such as parametric variation concerning discount rates. Consequently, the distribution of SCCs in the literature may not reflect experts' overall views about appropriate values for the SCC, *all things considered*.

If you agree to take part, you will first be asked to estimate the distribution of SCC values in the literature. You will also be asked to evaluate the effects on the SCC of key structural changes to IAMs. Finally, you will be asked for your assessment of the "true" distribution of the SCC, *all things considered*, as well as its drivers.

Taking part in this research study is completely voluntary.

You are free to decline to take part in this project. Please try to provide complete responses. However, you can decline to answer any question and you can stop taking part in the survey at any time, i.e., we also accept partial responses. Additionally, we encourage you to contextualize your quantitative responses with qualitative remarks.

You can also provide us with your name. Your name would allow us to check for non-response biases, to prevent multiple participations and to link responses to observable characteristics, such as relating to the SCC estimates from your published paper(s). If you choose not to, your answers will be anonymous. In any case, all results will be published in a way that no individual participant can be identified.

Questions

If you have any questions about this research, please feel free to contact the investigators, Moritz Drupp (Moritz.Drupp@uni-hamburg.de or +49-151-21221557) or Frances Moore (fmoore@ucdavis.edu or +1-617-233-3380).

Start the Survey

Many thanks for lending your expertise to this research!

Please click on "Next" to start the survey.

Next

Figure S10: Screenshot of the Preamble.

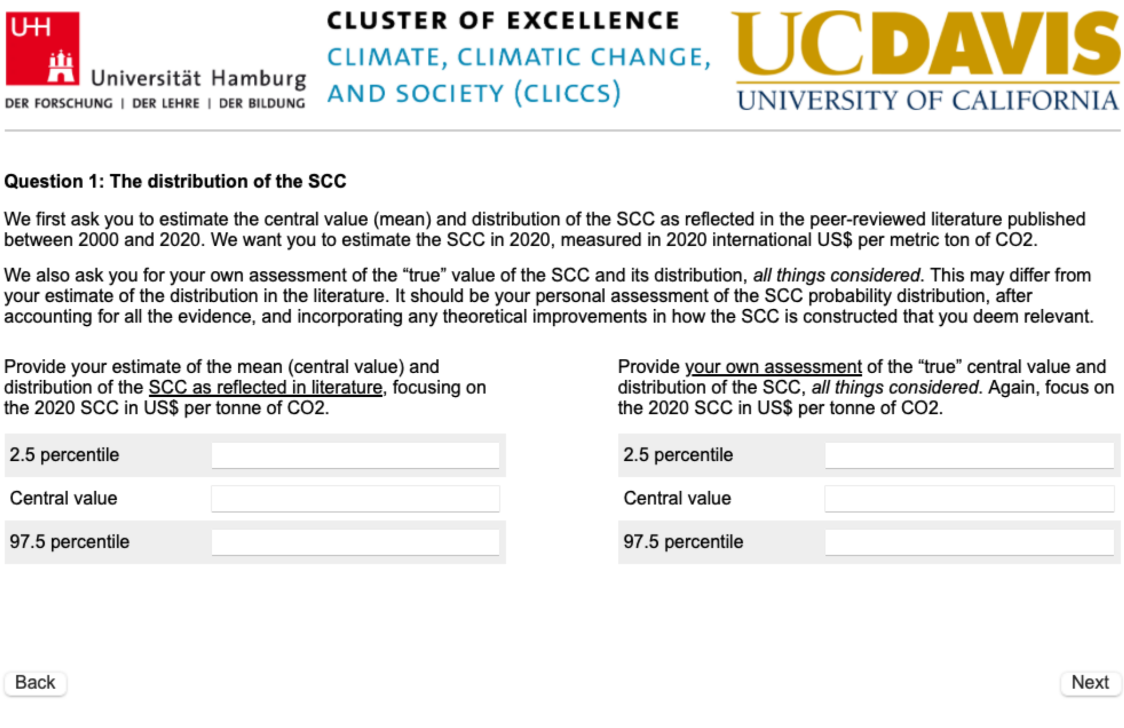

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CLUSTER OF EXCELLENCE
CLIMATE, CLIMATIC CHANGE, AND SOCIETY (CLICCS)

UCDAVIS
UNIVERSITY OF CALIFORNIA

Question 1: The distribution of the SCC

We first ask you to estimate the central value (mean) and distribution of the SCC as reflected in the peer-reviewed literature published between 2000 and 2020. We want you to estimate the SCC in 2020, measured in 2020 international US$ per metric ton of CO2.

We also ask you for your own assessment of the "true" value of the SCC and its distribution, *all things considered*. This may differ from your estimate of the distribution in the literature. It should be your personal assessment of the SCC probability distribution, after accounting for all the evidence, and incorporating any theoretical improvements in how the SCC is constructed that you deem relevant.

Provide your estimate of the mean (central value) and distribution of the SCC as reflected in literature, focusing on the 2020 SCC in US$ per tonne of CO2.

2.5 percentile	
Central value	
97.5 percentile	

Provide your own assessment of the "true" central value and distribution of the SCC, *all things considered*. Again, focus on the 2020 SCC in US$ per tonne of CO2.

2.5 percentile	
Central value	
97.5 percentile	

Back

Next

Figure S11: Screenshot of Question 1.

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CLUSTER OF EXCELLENCE
CLIMATE, CLIMATIC CHANGE, AND SOCIETY (CLICCS)

UC DAVIS
UNIVERSITY OF CALIFORNIA

Question 2: Structural drivers of the SCC in the literature

We now ask you about the effect on the SCC of *structural changes*, which have been considered in the literature between 2000 and 2020. Note that the structural changes appear in random order, i.e., nothing is implied about their relative importance.

To fix ideas and ease comparability across respondents, we ask you to consider a mean SCC of around \$40/tCO2 in 2020 international US\$ as the baseline or reference case *without structural changes*, i.e. not including any of the model characteristics listed below, and formulate effect sizes relative to this reference. This SCC value is, for instance, close to the central 2020 SCC of \$41.50/tCO2 in Bill Nordhaus' DICE2016R2 model (Nordhaus, 2018, *American Economic Journal: Economic Policy*).

We also invite you to comment on or further explain your answers.

Including tipping points in the climate system (e.g., Amazon forest dieback, ocean methane hydrate release)

a. On average, across papers that make these changes in the literature, how does this structural change affect the mean SCC?

○	○	○	○	○	○	○	○	○	○	○	○	○
< -200%	- 200% to -100%	-100% to -50%	-50% to -30%	-30% to -10%	-10% to 0%	0%	0% to 10%	10% to 30%	30% to 50%	50% to 100%	100% to 200%	> 200%

b. Do you agree with the statement that "papers in the current literature that incorporate this structural change produce a better estimate of the SCC than papers that exclude it"?

○ Strongly Disagree ○ Disagree ○ Neither Agree nor Disagree ○ Agree ○ Strongly Agree

Allowing for persistent effects of temperature change on output (e.g., via effects on the capital stock or TFP growth rate)

a. On average, across papers that make these changes in the literature, how does this structural change affect the mean SCC?

○	○	○	○	○	○	○	○	○	○	○	○	○
< -200%	- 200% to -100%	-100% to -50%	-50% to -30%	-30% to -10%	-10% to 0%	0%	0% to 10%	10% to 30%	30% to 50%	50% to 100%	100% to 200%	> 200%

b. Do you agree with the statement that "papers in the current literature that incorporate this structural change produce a better estimate of the SCC than papers that exclude it"?

○ Strongly Disagree ○ Disagree ○ Neither Agree nor Disagree ○ Agree ○ Strongly Agree

Figure S12: Screenshot of Question 2, Part 1.

Incorporating aversion to model uncertainty or ambiguity (e.g., by modelling second-order probabilities)

a. On average, across papers that make these changes in the literature, how does this structural change affect the mean SCC?

○	○	○	○	○	○	○	○	○	○	○	○	○
< -200%	- 200% to -100%	-100% to -50%	-50% to -30%	-30% to -10%	-10% to 0%	0%	0% to 10%	10% to 30%	30% to 50%	50% to 100%	100% to 200%	> 200%

b. Do you agree with the statement that "papers in the current literature that incorporate this structural change produce a better estimate of the SCC than papers that exclude it"?

○ Strongly Disagree ○ Disagree ○ Neither Agree nor Disagree ○ Agree ○ Strongly Agree

Using distributional or equity weighting

a. On average, across papers that make these changes in the literature, how does this structural change affect the mean SCC?

○	○	○	○	○	○	○	○	○	○	○	○	○
< -200%	- 200% to -100%	-100% to -50%	-50% to -30%	-30% to -10%	-10% to 0%	0%	0% to 10%	10% to 30%	30% to 50%	50% to 100%	100% to 200%	> 200%

b. Do you agree with the statement that "papers in the current literature that incorporate this structural change produce a better estimate of the SCC than papers that exclude it"?

○ Strongly Disagree ○ Disagree ○ Neither Agree nor Disagree ○ Agree ○ Strongly Agree

Structural changes to the temperature response to emissions (i.e., to the earth system module(s))

a. On average, across papers that make these changes in the literature, how does this structural change affect the mean SCC?

○	○	○	○	○	○	○	○	○	○	○	○	○
< -200%	- 200% to -100%	-100% to -50%	-50% to -30%	-30% to -10%	-10% to 0%	0%	0% to 10%	10% to 30%	30% to 50%	50% to 100%	100% to 200%	> 200%

b. Do you agree with the statement that "papers in the current literature that incorporate this structural change produce a better estimate of the SCC than papers that exclude it"?

○ Strongly Disagree ○ Disagree ○ Neither Agree nor Disagree ○ Agree ○ Strongly Agree

Including tipping points in the damage function(s)

a. On average, across papers that make these changes in the literature, how does this structural change affect the mean SCC?

○	○	○	○	○	○	○	○	○	○	○	○	○
< -200%	- 200% to -100%	-100% to -50%	-50% to -30%	-30% to -10%	-10% to 0%	0%	0% to 10%	10% to 30%	30% to 50%	50% to 100%	100% to 200%	> 200%

Figure S13: Screenshot of Question 2, Part 2.

Incorporating Epstein-Zin preferences (to allow for differentiated time and risk preferences)

a. On average, across papers that make these changes in the literature, how does this structural change affect the mean SCC?

< -200%	- 200% to -100%	-100% to -50%	-50% to -30%	-30% to -10%	-10% to 0%	0%	0% to 10%	10% to 30%	30% to 50%	50% to 100%	100% to 200%	> 200%

b. Do you agree with the statement that "papers in the current literature that incorporate this structural change produce a better estimate of the SCC than papers that exclude it"?

Strongly Disagree | Disagree | Neither Agree nor Disagree | Agree | Strongly Agree

Allowing for learning (e.g., about damages or tipping points)

a. On average, across papers that make these changes in the literature, how does this structural change affect the mean SCC?

< -200%	- 200% to -100%	-100% to -50%	-50% to -30%	-30% to -10%	-10% to 0%	0%	0% to 10%	10% to 30%	30% to 50%	50% to 100%	100% to 200%	> 200%

b. Do you agree with the statement that "papers in the current literature that incorporate this structural change produce a better estimate of the SCC than papers that exclude it"?

Strongly Disagree | Disagree | Neither Agree nor Disagree | Agree | Strongly Agree

Comments

Back | Next

Next

Figure S14: Screenshot of Question 2, Part 3.

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CLUSTER OF EXCELLENCE
CLIMATE, CLIMATIC CHANGE, AND SOCIETY (CLICCS)

UCDAVIS
UNIVERSITY OF CALIFORNIA

As you did not provide numerical central values for the SCC, for completing the next question we now need to know whether your own assessment of the "true" central value of the SCC, *all things considered*, is higher, lower, or the same as the central SCC estimate in the literature?

"True" SCC is higher | "True" SCC is lower | "True" SCC is the same

Back | Next

Figure S15: Screenshot of Interim Question in case the mean estimates for the SCC from Q1 were not answered (completely).

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CLUSTER OF EXCELLENCE
CLIMATE, CLIMATIC CHANGE, AND SOCIETY (CLICCS)

UC DAVIS
UNIVERSITY OF CALIFORNIA

Question 3: Drivers of the SCC wedge

We now ask you what drives any difference between (1) your own assessment of the "true" central value of the SCC, *all things considered* and (2) your estimate of the central value of the SCC in the literature. We call this the 'SCC wedge'.

Below, we provide a list of potential drivers of the SCC wedge, and we ask you to weigh their importance. Each individual weight can be in the range +100% to -100%. The default when not clicking an option is 0, i.e. that this driver does not affect your SCC wedge.

The list below includes structural and parametric variations. It is again in random order. It is not exhaustive, so we provide a residual, *Other drivers* category for any remaining drivers you think should be included. Please provide bullet points on which additional drivers you have considered in the text box below.

As your *all things considered* SCC is higher than your estimate of the SCC in the literature, and your SCC wedge is positive, you should ensure that all weights you put on the drivers plus the "Other drivers" category add up to +100 (in %).

In the box below, you see the remaining budget of weights to be allocated. You will also get a notice of the remaining budget to allocate if you click on "Next" at the bottom, with the option to adjust, rescale or skip this question.

Drivers of the SCC wedge:

What drives the difference between your *all things considered* central value for the SCC and the central SCC value you perceive in the literature?

Remaining weights to be allocated: 100

	-100	0	100
Adjusting pure time preference		0	
Allowing for (endogenous) adaptation to climate damages		0	
Allowing for persistent effects of temperature change on output		0	
Incorporating Epstein-Zin preferences		0	
Incorporating limited substitutability between market and non-market goods		0	
Updating the earth system module(s)		0	

Figure S16: Screenshot of Question 3, Part 1.

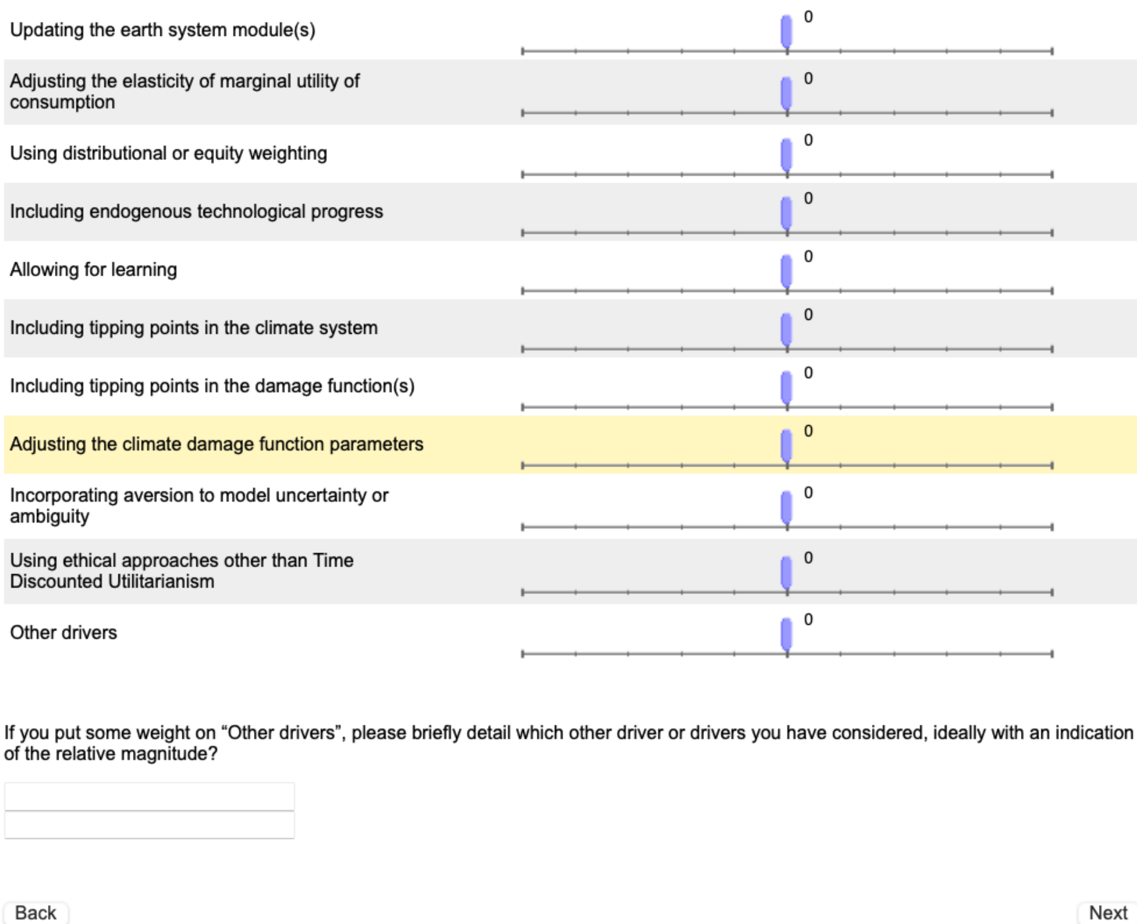

Updating the earth system module(s)	0
Adjusting the elasticity of marginal utility of consumption	0
Using distributional or equity weighting	0
Including endogenous technological progress	0
Allowing for learning	0
Including tipping points in the climate system	0
Including tipping points in the damage function(s)	0
Adjusting the climate damage function parameters	0
Incorporating aversion to model uncertainty or ambiguity	0
Using ethical approaches other than Time Discounted Utilitarianism	0
Other drivers	0

If you put some weight on “Other drivers”, please briefly detail which other driver or drivers you have considered, ideally with an indication of the relative magnitude?

Back

Next

Figure S17: Screenshot of Question 3, Part 2.

CLUSTER OF EXCELLENCE
CLIMATE, CLIMATIC CHANGE, AND SOCIETY (CLICCS)

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

UC DAVIS
UNIVERSITY OF CALIFORNIA

Question 4: Comments on next steps in improving the SCC

What do you think are the most important steps for improving estimates of the SCC going forward?

Your name

Finally, may we kindly ask you to provide your name? Your name would allow us to check for non-response biases, to prevent multiple participations and to link responses to observable characteristics, such as relating to the SCC estimates from your published papers.

We will hold your responses in the strictest confidence.

Upon closing the survey, we will create a separate password secured ID file that links identities to ID numbers and perform matching to meta-analysis data on this basis. The link of IDs to identities will be stored in password secured file to which only Moritz Drupp and Frances Moore have access. Also, the full dataset will be password secured. Results will only be published anonymously and such that no individual participant can be identified. Specifically, we will publish the meta-analysis data separately from the survey data, to ensure that no individual can be re-identified, and will only share data at an aggregated level with our co-authors on the project (Simon Dietz, James Rising, Ivan Rudik, Gernot Wagner).

Family name	
First name(s)	

We naturally also welcome anonymous responses if you do not want to reveal your identity.

Additional feedback

Feel free to provide us with any additional comments or feedback:

Many thanks for your valuable contribution!

Best regards,
Moritz Drupp (Hamburg) and Frances Moore (UC Davis)

Back | Next

Figure S18: Screenshot of Question 4 and Ending.

S.2.2.2 Survey response cleaning

We took the following steps to clean the response data, indicated by row values #XXX, for final analysis. Row values of 1000 and beyond relate to response received via e-mail. Specifically, we

- removed one duplicate response, which was a subset (#698, retained #699)
- changed 2.5 and 9.75 percentiles for a single case in which they were entered reversely (#718)
- disregarded responses to Q3 of respondents who did not move any of the cursors.
- disregarded responses to Q3 in one case that only provided an all-things-considered SCC but no literature value and where the weights did not add up (#682)
- disregarded responses to Q3 in one case where weights did not add up and respondent ticked "Do not answer this question." (#626)
- disregarded responses to Q3 in one case where weights did not add up, appeared extremely strange and the respondent was anonymous so that it was not possible to follow-up.
- re-weighted weights in Q3 to add up for one respondent (#599) whose weights did not add up and who ticked "Rescale my weights such that they add up." in the survey
- added additional response categories, such as explained non-responses for authors who had retired in the meantime
- corrected SCC values of one expert (#576) based on e-mail communication following-up on a comment in the survey (see below) that the respondent had missed the option to go back within the survey to adjust the SCC estimates for equity weighting.
- de-anonymized one expert after the respondent had identified themselves and their response bilaterally (#750)
- added qualitative responses to Q4 from four experts who respondent via e-mail, and a quantitative response to the SCC wedge for one respondent (#1000).

S.2.2.3 Publication bias analysis

Figure 1 shows the distribution of SCCs published in peer-reviewed journals over the 2000-2020 period. It is important that this is not interpreted as a standard meta-analysis. In the classic meta-analysis, multiple studies have produced empirical measurements of the same quantity, which can be combined to give lower-variance estimates of the quantity of interest. The SCC estimates we bring together here are fundamentally distinct in that they are not observational measurements of an empirical quantity, but primarily results of model simulations. The variation in Figure 1 reflects not uncertainty related to statistical sampling (as in a classic meta-analysis), but epistemological uncertainty in model parameters and structure. The distribution should be thought of as an integration over this uncertainty as reflected in the published literature over the last 20 years.

This interpretation does raise the question of how different the distribution would be if it also included unpublished SCC estimates (or those published but not in peer-reviewed journals). The question of publication bias arises repeatedly in the context of standard meta-analyses, where it refers to missing evidence of small effect sizes due to a lack of incentives to publish null effects. It is not clear whether similar asymmetric publication incentives operate around SCC values and, if so, in what direction they would shift the distribution. On the one hand, there is some evidence of a conservatism in scientific publishing, such that one might expect an anchoring around previously published estimates and therefore a narrower distribution in the published literature compared to unpublished model results (*20, 21*). On the other hand, others might suggest that extreme SCC values (either very high or very low) might be more noteworthy and so be more likely to proceed to publication, implying a wider distribution in the published literature. Existing studies show evidence in both directions, with Havranek et al. (*22*) finding substantively lower average SCC values in peer-reviewed published journals than those published in other outlets while a review by Tol (*23*) finds the opposite.

Our expert survey was designed to fill key knowledge gaps on the SCC that the meta-analysis alone cannot answer: the role of publication bias and insights on key drivers of the SCC and next steps for improving its estimation. The standard concern regarding publication bias is that due to researcher and editorial incentives, leading to file drawer problems or questionable to fraudulent research practices (e.g. p-hacking, data fabrication), estimates of a true effect size is represented in a biased form in the literature. Standard approaches are to examine z-scores of irregularities in p-values for experimental studies that aim to investigate some true effect size. Yet, these are not directly applicable to our setting, as many individual SCC papers do not have the goal of producing the best estimate of the SCC. These papers are comparative in nature and explore the effect of some variation in plausible parameters or some extension of the IAM structure to investigate how this affects the SCC. Oftentimes, they do so starting from a well-established, conventional baseline, such as the latest DICE model, and consider only one or a few extensions. Thus, these studies will—by design—provide a 'biased' estimate of what the study authors may themselves consider an appropriate estimate of the "true" SCC.[1]

Previous work has investigated publication bias by comparing peer-reviewed and non peer-reviewed work on the SCC (*22*). Yet, for the reasons detailed above, comparing published and unpublished papers trying to estimate the SCC does not provide the most useful benchmark for detecting publication bias. We therefore investigate the role of publication bias by means of the expert survey. Specifically, any differences between an expert's estimate of the 2020 SCC in the literature and the true comprehensive value of the SCC serves as an indication of publication bias.

Our data suggests that, at least in the view of our surveyed experts, the peer-reviewed literature exhibits a substantial and significant downward bias in the central SCC as compared to experts' subjective "true" central estimate (t-test, t=6.063, two-sided P<0.000). This is also apparent given the substantial rightward shift in the "true" distribution—one that asked experts to account for all things they deemed missing or imbalanced in the published literature—compared to the estimated literature distribution in Figure 1. Indeed, we find that 82.82% of experts think the central SCC in the literature to be biased downwards as compared to its true value, 9.09% think that it is represented correctly in the literature and 9.09% think the SCC to be biased upwards in the literature.

[1] There may be an incentive—akin to p-hacking in the empirical literature—for such comparative studies, as the chance that they get published well may increase if the paper reports a strong effect of introducing a certain structural model modifications. The magnitudes of structural model modifications may thus be overestimated in the literature.

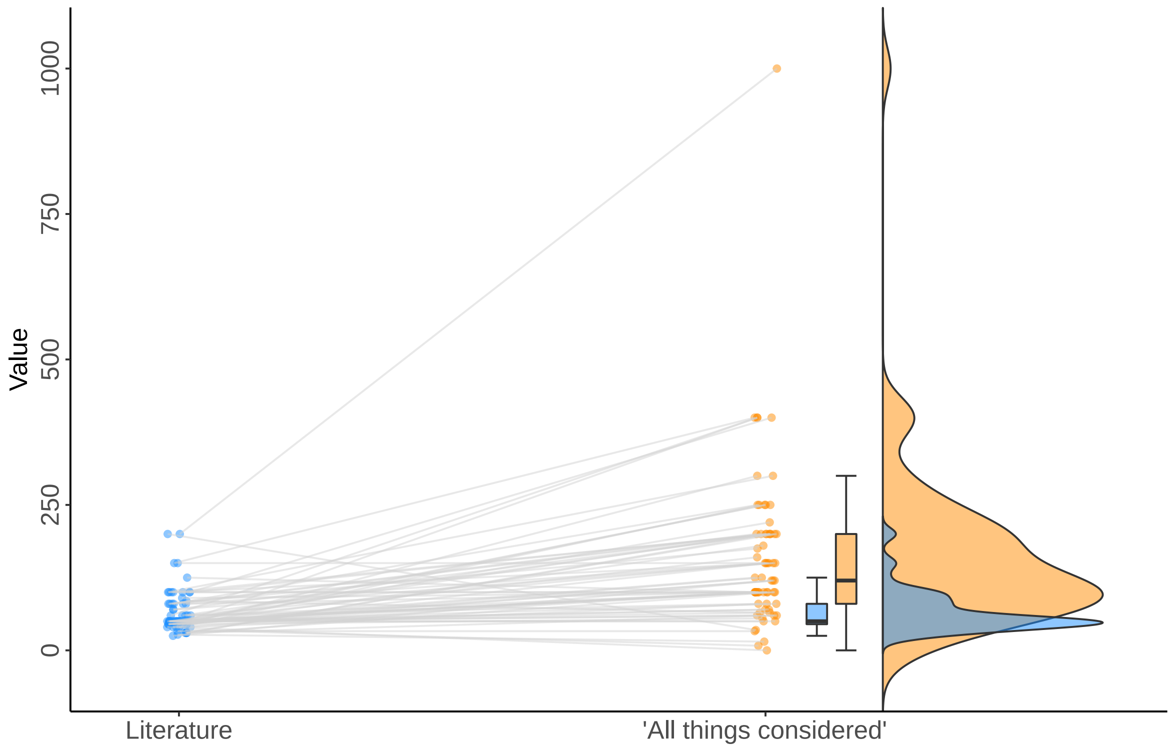

Figure S19: Individual estimates of the central 2020 SCC in the literature and the central "true" 2020 SCC value "all things considered".

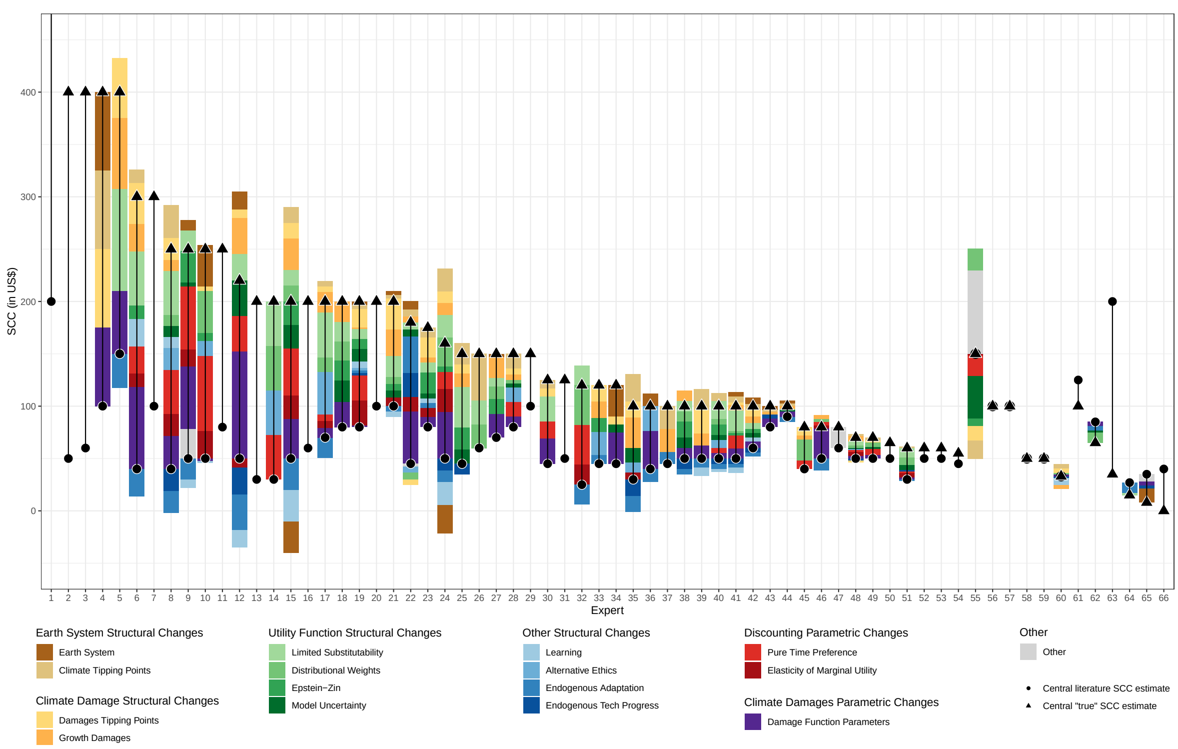

Figure S20: Individual estimates of the central 2020 SCC in the literature and the central "true" 2020 SCC value "all things considered".

S.2.2.4 Non-response and strategic response bias analysis

We conduct a number of non-response and strategic response bias checks (*15, 16*). To investigate non-response bias, we compare non-anonymous respondents with the rest of the population of authors along observable characteristics such as their continental location, gender and year of PhD, number of publications in SCOPUS and their h-index. We complement this with data from our meta-analysis on the median 2010-2030 SCC in their published papers, the median synthetic discount rate employed in their SCC papers, as well as the proportion of their SCC estimates that contain one of the major structural model modifications, and the proportion of their SCC estimates that we classified as 'framework expansion' or as 'empirical improvement'. Overall, we find that identified respondents ($n_{ir} = 48$) exhibit very similar characteristics as the anonymous part of the population sample ($n_{non-ir} = 128$), which contains the anonymous respondents ($n_{ar} = 27$). We find no significant differences across groups in terms of being located in North-America (t-test, t=0.970, two-sided P=0.334), Europe (t-test, t=-1.476, two-sided P=0.142), Asia (t-test, t=0.522, two-sided P=0.602) or Oceania (t-test, t=0.598, two-sided P=0.550), being classified as 'male' ($n_{ir} = 44$, $n_{non-ir} = 109$; $\chi^2(1)$=1.303; P=0.339), year of PhD award (t-test, t= -0.4799, two-sided P=0.632), number of publications (t-test, t=0.701, two-sided P=0.484) and h-index (t-test, t=0.570, two-sided P=0.0.569. We also find no significant differences in terms of their synthetic discount rates employed (t-test, t=-1.461, two-sided P=0.146) and their 2010-2030 SCC estimates (t-test, t=-0.189, two-sided P=0.851). Furthermore, identified respondents have published more frequently on what we classified as 'empirical improvement' (t-test, t=-2.432, two-sided P=0.016), they do not exhibit a higher proportion of having published a 'framework expansion' (t-test, t=-0.384, two-sided P=0.702), nor on any of the sub-classifications, such as on climate or damage tipping points (t-test, t=0.615, two-sided P=0.975; t-test, t=-1.461, two-sided P=0.540), earth system model updates (t-test, t=-1.281, two-sided P=0.202), alternative utility function specifications, e.g. using Epstein-Zin (t-test, t=-1.281, two-sided P=0.0.885), or alternative ethical approaches (t-test, t=0.283, two-sided P=0.778). This null finding is noteworthy as having mentioned some of these in the initial e-mail invitation might have induced experts who have published more on these to be more likely to respond to the survey.

To investigate strategic response bias, we compare responses by anonymous and identified respondents. We find no significant difference in views on the central "true" SCC all things considered (t-test, t=-0.991, two-sided P=0.325), and the upper and lower percentile ranges are almost identical (two-sided P=0.898 and P=0.949). While anonymous respondents estimate a much lower central 2020 SCC in the literature (t-test, t=-2.497, two-sided P=0.015), the upper and lower SCC literature range estimates do not differ (two-sided P=0.420 and P=0.241). Both groups also do not differ in terms of how frequently they think that the "true" central SCC is higher than the central SCC in the literature ($n_{ar} = 25$, $n_{ir} = 41$; $\chi^2(1)$=0.916; P=0.339), and in the quantitative size of their SCC wedges (t-test; t=-0.486, two-sided P=0.628). Further, we do not find any differences in their views on whether any of the structural model modifications represent an improved estimate of the SCC (t-tests, lowest P=0.140, for persistent growth effects), and anonymous respondents are also not more likely to provide comments on any of the questions (lowest $\chi^2(1)$=0.012 is P=0.358 for Q3). Yet, anonymous respondents put considerably more weight on 'other drivers' in question 3 (t-test, t=-3.027, two-sided P=0.004), and less weight on pure time discounting (t-test, t=1.731, two-sided P=0.090). Whereas anonymous respondents are not more likely to provide comments on structural model modifications in Q2 or on next steps for improving the SCC ($n_{ar} = 27$, $n_{ir} = 48$; $\chi^2(1)$=0.012; P=0.913). While responses of the anonymous sample are different along some dimensions, we do not detect a clear signal of efforts to strategically distort the survey results, in particular as the "true" comprehensive SCC does not differ across samples. Differences appear to rather stem from different views on the SCC in the literature and on the role of pure time preference in driving the SCC wedge. To be inclusive of the whole range of reasonable views, we thus retain anonymous responses as part of our main analysis. We repeat the strategic response bias analysis by splitting the sample at the median into early and late respondents. The general hypothesis is that respondents who want to strategically affect the results may respond earlier. We find no significant differences along all those dimensions reported on above, except that early respondents are more likely to provide comments on next steps for improving the SCC ($n_{early} = 33$, $n_{late} = 33$; $\chi^2(1)$= 7.174; P=0.007).

S.2.2.5 Additional analyses of the merged dataset

Among the identified respondents, a higher 2010-2030 SCC median in an expert's own publications is highly significantly associated with estimating a higher SCC in the literature (β=0.05, $t_3$3=4.50, P <0.000), and weakly significantly associated with a higher "true" SCC estimate (β=0.12, $t_3$3=1.75, P=0.089).

We do not find significant differences in literature or "true" SCC estimates across continents (but those by experts from Asia tend to be lower and those by experts from Europe tend to be higher).

We find a strongly significant gender effect: The few female experts estimate a significantly higher SCC value in the literature (t-test, t=4.269, two-sided P<0.001) and also a "true" SCC value that is more than twice as high as the corresponding estimate by male experts, with \$288 versus \$130 (t-test, t= 3.80, two-sided P<0.001), and a higher SCC-wedge (t-test, t= 2.74, two-sided P=0.009).

We next relate the proportion of an expert's included papers published on a specific structural model modification and whether they are more likely to agree or strongly agree that "papers in the current literature that incorporate this structural change produce a better estimate of the SCC than papers that exclude it" (Q2). We do not find this to be the case for any of the extensions. However, we find that having published (a higher share) of papers on some specific extensions—persistent growth damages (β=14.01, t_{40}=2.13, P=0.040), Epstein-Zin preferences (β=22.73, t_{40}=6.63, P <0.000), and distributional weights (β=27.34, t_{40}=2.87, P=0.007)—is associated with assigning more weight on these structural model modifications as drivers of the central SCC-wedge.

S.2.2.6 Survey Distribution

Experts were asked their estimates of the mean SCC (for both the literature and their comprehensive estimate) and the 2.5th and 97.5th percentile. We a similar algorithm as that described in S.2.1.4 for the literature distribution to find best-fit distributions for the reported values, though with a focus on matching the surveyed mean value over the surveyed tails, following our principle above to assume short tails when possible.

The fitting algorithm proceeds as follows:

If the expert only provided a mean estimate without quantiles, the SCC is assumed to be deterministic at that central value.

If the expert provided one quantile (2.5% or 97.5%), a symmetric triangular distribution is assumed.

If the expert provided two quantiles (2.5% or 97.5%), a piecewise uniform distributions is assumed with two segments. The mid-point and width of these two distributions is uniquely determined by the quantiles and mean value. We further assume symmetric tails, using the tail fitting method above to generate left and right tails and then treating each tail as a mixture of these two results. In some cases (three total responses), no distribution following these assumptions can fit the reported data, and these entries are dropped from the combined distribution.

Expert distributions shown in Figure 1 are based on 1000 samples from each fitted distribution from each expert. This therefore incorporates elicited tail information from experts in addition to the central estimates. The results are reported in Table S8. The mean of \$158 differs slightly from the mean of reported mean values taken directly from the survey, due to dropping three responses for which distributions could not be identified.

	2.5%	5%	25%	50%	75%	95%	97.5%	Mean
Literature	-4	3	18	34	62	245	450	65
Comprehensive	-7	3	41	85	166	581	1055	158

Table S8: **Fitted expert survey SCC estimates.** We elicited the central (mean) SCC and the 2.5 and 97.5 ranges. From this, we fit a istribution of expert survey estimates of the 2020 SCC in the literature and the comprehensive SCC, all things considered. Additional details in Methods.

S.2.2.7 Survey Meta-analysis

We apply a form of Bayesian hierarchical modeling to interpret survey responses to the question "To what extent do you agree with the statement that papers including X produce a better SCC than those excluding it?" The model makes the following assumptions:

- Each expert has a belief about the likelihood that including a given structural model modification is beneficial. This believe is uncertain, and represents a distribution over probabilities.
- The response categories– "Strongly Disagree", "Disagree", "Neither Agree nor Disagree", "Agree" and "Strongly Agree"– represent a discretization of this probability space. To choose one category, an expert will take a random draw from their belief distribution.
- Each expert has a consistent ordered scheme for mapping probability values to categories. Also, they accurately report their beliefs.
- We can usefully talk about the "common belief" across experts. This is a (hyper-)distribution, from which each expert's beliefs are drawn. It is possible to partially pool the beliefs across the experts to generate an estimate of this common belief.

This represents a kind of meta-analysis, which allows us to simultaneously generate an estimate of the common beliefs about structural model modifications, and expert beliefs that are consistent with these. If the expert beliefs are estimated to be very uncertain, a high level of pooling will be used since all estimates will be consistent with a common value. If expert beliefs are certain and not alike, little pooling will be used and the estimate of the common belief will be uncertain.

We use the following Bayesian model to implement these ideas. Let the central logit value of the belief of expert j on question k be

$$\theta_{jk} \sim \mathcal{N}(\mu_k, \tau_k)$$

The actual reported level is a categorical variable drawn from an ordered logistic:

$$l_{jk} \sim \text{OrderedLogistic}(h\theta_{jk}, h\vec{c}_j)$$

where h is a global parameter that determines the spread around the central value (that is, how likely it is for an expert to report a category higher or lower than their central belief).

The categorical divisions are deviations from a prior understanding of the categories, represented by

$$c_j \sim \mathcal{N}([\text{logit}(0.2), \text{logit}(0.4), \text{logit}(0.6), \text{logit}(0.8), \text{logit}(1.0)], \sigma)$$

where σ determines the variation between expert understandings of the categories.

We further impose weakly informative priors that $h \sim \text{Exponential}(0.001)$ and $\tau_k \sim \text{Cauchy}(0, 1)$.

The result of this analysis is shown in Figure S21. Some questions show a higher degree of agreement (e.g., ambiguity/model uncertainty) than others (e.g., Epstein-Zin preferences). The mean of the common belief estimate is always very close to a simple average of experts' answers. However, the range of probabilities for the common belief, with a standard deviation of about 0.1, is considerably better constrained than the original values. Summary statistics and parameter values are shown in Table S9.

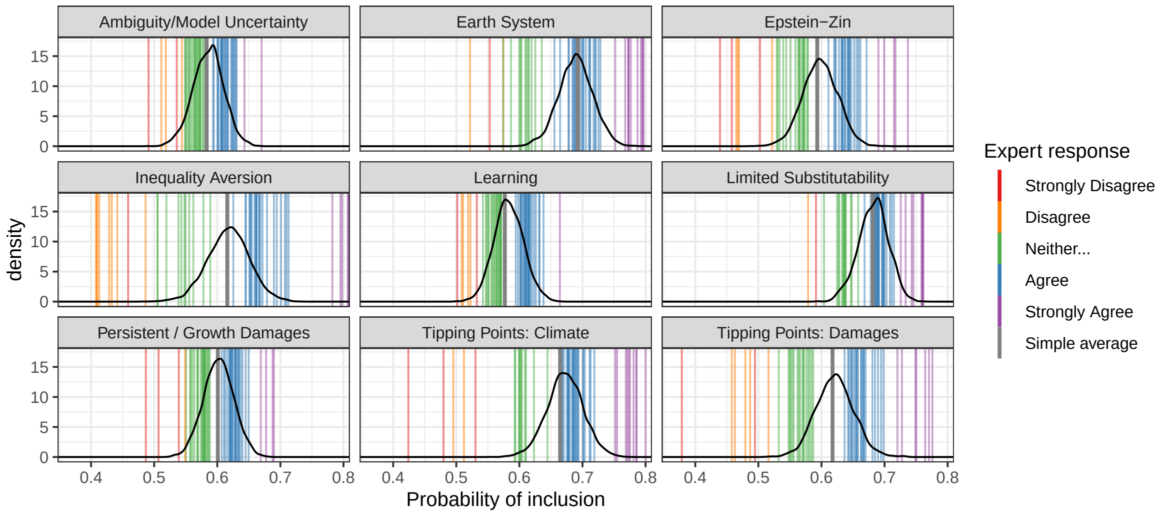

Figure S21: The Bayesian hierarchical modeling estimates of expert central beliefs. Each expert is shown as a line, colored by the category of their answer and placed along the x-axis at their estimate belief level. The distribution of the common belief is shown in black, and a simple average of expert opinions, assuming that reported answers represent the center of five probability bounds, is shown in grey.

	mean	se_mean	sd	2.5%	25%	50%	75%	97.5%	n_eff	Rhat
Ambiguity μ	0.35	0.00	0.10	0.15	0.28	0.35	0.42	0.56	5881.72	1.00
Earth System μ	0.79	0.00	0.12	0.56	0.71	0.79	0.87	1.05	4277.46	1.00
Epstein-Zin μ	0.40	0.00	0.12	0.16	0.32	0.40	0.48	0.63	4430.73	1.00
Inequality μ	0.48	0.00	0.14	0.21	0.39	0.49	0.58	0.76	3132.52	1.00
Learning μ	0.34	0.00	0.10	0.15	0.27	0.34	0.41	0.54	7069.27	1.00
Limited Sub. μ	0.77	0.00	0.11	0.56	0.70	0.77	0.85	1.00	5690.99	1.00
Persistence μ	0.42	0.00	0.10	0.22	0.34	0.41	0.49	0.62	5665.33	1.00
TPs: Climate μ	0.73	0.00	0.13	0.47	0.64	0.73	0.81	0.98	4692.44	1.00
TPs: Damages μ	0.49	0.00	0.12	0.25	0.41	0.49	0.58	0.74	4303.60	1.00
Ambiguity τ	0.19	0.01	0.14	0.01	0.08	0.17	0.28	0.51	609.14	1.00
Earth System τ	0.44	0.01	0.19	0.04	0.31	0.45	0.57	0.81	503.55	1.01
Epstein-Zin τ	0.39	0.01	0.19	0.03	0.25	0.39	0.52	0.76	443.44	1.01
Inequality τ	0.68	0.01	0.19	0.25	0.58	0.69	0.80	1.04	249.03	1.02
Learning τ	0.21	0.01	0.14	0.01	0.10	0.20	0.31	0.51	577.08	1.01
Limited Sub. τ	0.29	0.01	0.18	0.02	0.15	0.28	0.42	0.65	598.69	1.01
Persistence τ	0.22	0.01	0.16	0.01	0.09	0.20	0.33	0.56	506.57	1.00
TPs: Climate τ	0.48	0.01	0.22	0.04	0.33	0.49	0.64	0.90	364.83	1.01
TPs: Damages τ	0.49	0.01	0.20	0.07	0.36	0.51	0.63	0.86	430.27	1.01
h	3.22	0.05	0.54	2.57	2.89	3.12	3.40	4.60	130.10	1.03
σ	0.36	0.00	0.04	0.29	0.34	0.36	0.39	0.45	799.47	1.00

Table S9: Posterior distribution statistics for common parameters, computed by Stan, using 4 chains, each with iter=2000; warmup=1000. n_eff is a measure of the effective number of MCMC draws of the posterior distribution that were achieved. Rhat is a measure of convergence, where full convergence produces a value of 1.0.

S.2.2.8 Qualitative comments

Below we report qualitative comments received in the online survey or via e-mail, not edited except in cases where this is necessary to preserve anonymity. Numbers refer to individual experts by row values (#XXX). Row values of 1000 and beyond relate to response received via e-mail.

Comments on Question 2: Structural drivers of the SCC in the literature

#553 Persistent effects are still under debate about how to integrate them in IAM. Current implementations are not satisfying.

#558 My responses to b are all else equal - I would have preferred this to say the same model or paper with or without a specific feature.

#562 Uncertainty would not alter SCC that much. Exceptions would include uncertainty about the damage function related to tippiing points or uncertainty about how rapid the temperature increases I believe.

#569 For most of these I put "Neither agree or disagree" not because I don't think they are important to develop further, but because I don't think the evidence base is quite there yet or because I think there are a set of difficult ethical assumptions inherent in making these changes that I'm not sure are better than the current assumptions.

#572 most of my responses are very impressionistic based on recollections of a partial and idiosyncratic reading of the literature. would have to conduct a more systematic review to have any confidence in my answers. but i assume what you want here are impressions

#576 When producing my own distribution in the previous page, I really focused on damages and discount rate holding structure constant. However, I would note that some of these parameter changes do not impact the social cost of carbon as much in my opinion, but instead the optimal tax. This is particularly true for learning by doing, as learning by doing can substantially impact the optimal tax, though I think it will have a limited effect on the most likely climate scenario.

#598 In simple models, a lot depends on how the damage function is calibrated. A simple damage function could also include damage on capital, tipping points, non-market goods, etc. In models with learning on damages, I assume that the model without learning does monte-carlo, does not have the option to change the plan when information is discovered and has therefore a much larger risk premium.

#623 I don't think "tipping points" as advanced by Lenton, Schellnhuber, etc are physically justified by the available science. I think the main thing missing in terms of losses associated with CC is a reasonable representation of the on-going costs of extreme events. These have been mostly neglected in the development of DICE-like IAMs. I think lots of ideas around equity-weighting in SCC estimations start to lose contact with political reality quite fast. I'm not exactly sure how to bound "structural changes" so I down-weighted the potential effects. These could be higher, given some possible interpretations.

#628 it is hard to distangle the impacts of uncertainty and learning in the questions. Generallyu I would think that uncertainty (for example in ECS, tipping points etc) increases SCC, but with that uncerrtainty eventual learning of the issue at hand (eg ECS) reduces the SCCncreases

#633 Each of these questions merits a proper meta survey paper so my answers are highly noisy and should not be taken too serious at all

#645 Some of the statements were ambiguous, like 'Structural changes to the temperature response to emissions', 'tipping points in the damage function' or 'persistent effects of temperature change on output'. It was very hard to evaluate what these actually mean, what their inclusion means (i.e. how are they included in a model) and thus what impact they could have on the SCC. Model uncertainty and ambiguity are very important concepts in this context. But are they something that we can objectively judge and quantify in the SCC? Or are they rather points that we should recurrently remind ourselves and reflect upon. Yes, there are mathematical theories that allow one to incorporate second-order probabilities and the aversion to ambiguity. But does that really capture all the ambiguity there is, or do these theories just convey 'artificial crispness' (as Weitzman described this) that is not truly the thing we are after?

#659 Re some of the above topics I am aware of only one study, in others not even one. So difficult to come up with my guesstimate.

#660 Why did you leave out CO2 fertilization effects? That's one of the big differences among IAMs and SCC estimates. I took the "structural changes to temperature response" to mean the Equilibrium Climate Sensititivity parameter. This is one of the big uncertainties in climate modeling not only among models but between model-based and empirical estimates. The term "tipping points" is overly vague. The proper term is "bifurcations". If they exist they are properties of the climate system, they are not induced by forcings. The literature "incorporating tipping points" gives a sense of looking for ways to make a simulation model crash and generate a dramatically larger SCC, without explaining how such a bifurcation could have existed throughout the present and previous interglacials without causing similar crashes even during much warmer periods. This genre looks to me to lead to higher SCC's of lower scientific quality. Adding growth impacts from temperature and precip changes in principle should increase the quality of the empirical estimate but in practice the results are very uncertain and of indeterminate sign. So they look to me like better quality but no change to the mean or median estimate. I don't have any familiarity with the techniques for incorporating aversion to ambiguity, Epstein-Zin preferences or constrained substitutability.

#687 On inequality, it depends on what you mean by inequality. Current estimates show convergence of incomes so this would increase the discount rate and reduce SCC. If you add in the inequality of the impacts of climate change,, which are borne by the poor, then this could well increase the SCC. If you then introduce catastrophic effects to the poor, you will have dismal theorem type effects which could be really important welfare wise. None of this nuance can be captured by the questions. On the climate science questions: whether updating the energy balance model will increase or decreas the SCC, I am afraid that I have not read so many papers on this topic, so i do not have a good basis for giving answers here. On tipping points, this is a dificult question. We know so little about how likely they are to occur and how bad tipping points will be in terms of economic damages that it is a difficul tcall to say that an estimate of the SCC for policy purposes is better for their inclusion. What is required is a presentation of the range of what is known, coupled with some decision / welfare theoretical approach to assess what the welfare maximisng approach is when faced with such dramatic uncertainties (a la Barro, Weitzman etc.). recent work by Dietz et al (2021) tends to indicate that tipping points will be trivial in terms of impact on GDP due to physial factors (e.g. feedbacks) but also economic factors (discounting, damage function and the fact that an optimal path will reduce the risks considerably) . So, difficult to say whether I think they should inform THE SCC,but they should inform the range of possibilities that are presented to policy makers for sure. the above comments are written rather quickly. I think they make sense. Happy to clarify ex post if need be.

#738 I found it really hard to answer these questions. First of all because I do not think there is a true SCC. The SCC depends on valye laden assumptions such as the discount rate, how to value the loss of a statistical life in poor countries, how to value risk, and how to value nature... hence, there can be no true answer here. Furthermore - how a change in the model will affect the SCC depends of course on how the change was implemented and how that structural phenomena was implemented before... so for some of these questions I was uncertain about how to ask. ,

#749 The only truly realistic change is including uncertainty in the scc rather than speaking of a distirbution of scc, and that change will raise the scc

#752 I find this page of questions somewhat ambiguous as the answer usually depends on the details of what is considered "structural change". Depending on these details in many if not most cases the answer to "part a" can change the sign and the answer to "part b" can change from Agree to Disagree. I decided in each case what I guessed you might be after, and I did not represent the literature per se but that part of the literature corresponding to my guess of what you might be after. E.g. growth change: even the DICE model has a pass through of the damages on the capital stock (so yes, agree, important), but I assumed you are likely after assuming direct damages on the growth rate rather than output or capital stock (then answer is, no, disagree that it improves the estimates conditional on current knowledge). There are similar ambiguities in most questions.

Comments on Question 3: Drivers of the SCC wedge.
"If you put some weight on "Other drivers", please briefly detail which other driver or drivers you have considered, ideally with an indication of the relative magnitude?"

#576 I did not consider equity weights in my $200. If I considered equity, then I would have multipled my estimate by 2.5 coming to an estimate of $500. I put 5% on the elasticity of marginal utility of consumption and Epstein Zinn preferences each, as these are really a joint decision in my opinion. I think though that market rates have declines, which further supports these differing parameter values underlying an extended Ramsey equation.

#617 Difficult for me to answer. My main thinking and own contribution have concerned the tail probabilities.

#626 hard to grasp this question I am afraid

#641 Shares with sliders?

#650 Generally, nonconstant discounting and monetizing more sources of damages should lead to a higher SCC. Adaptation and more granular (geographically) treatment of damages should decrease the SCC.

#660 Other drivers: I include endogenous adaptation here (I wasn't sure if that's what you meant by allowing for learning). Updating the earth system module: I take this to mean updating the ECS parameter to empirically-constrained ranges Adjusting the climate damage function: I include incorporation of CO2 fertilization and global greening effects here.

#679 Measurability problem in climate damage functions: SCC considers only measurable/marketable damages. F.e. costs of "wildlife extinction" not considered/not measurable in terms of GDP/welfare loss.

#687 I didn't say other but, my responses are based on the fact that one does not need too much in the way of other changes to get to the SCC that I would prefer. Many of the other points above are important, but not necessary to explain the difference. there may be some discrepancy between these responses and the previous responses, but i am explaining here the primary deerminants of the gap in terms of my original position on the SCC which rested on the above aspects.

#697 Impact on sentient animals

#738 Of course - this is tricky to answer... so please see my answers as rough estimates. For instance - changes in the damage function, could be done by change ion the parameters or by adding tipping points. It is not trivial to disentangle these.. the same with the discount rate and eta vs rho

#742 I will need to go back and do the first question. I think it would be better to have one think through the components first, no?

Question 4: Comments on next steps in improving the SCC.
"What do you think are the most important steps for improving estimates of the SCC going forward?"

#553 Understanding the impact of damage persistence and the interplay with adaptation.

#558 better incorporation of the latest science in the earth system modeling; ability to more adequately capture tipping points, including economic ones such as collapse of maritme food sources and massive migration better modeling of how economic systems and prices will adjust; better modeling of uncertainty

#562 SCC fully considering the effect of learning on multiple uncertain features at the same time

#567 Transparency about modelling choices and their effects on the SCC; improving climate damage estimations; explicity taking into account nature's values and relative price changes; societal discussions and refelctions about intergenerational equity and fairness

#569 Exploring tipping points, better characterizing damage functions, improving exploration of different ethical frameworks.

#572 improved evidence base for calibration of damage function parameters. i.e., more empirical studies of the influence of climate changes on economic performance and well-being.

#579 macro-economic modeling to understand pathways by which climate change might haver persistent effects on the economy

#598 Combine risk and non-marketed goods, a more dynamic damage function (warming speed matters and temperature may affect tfp growth) It is not widely known that in most models the risk premium is actually low (increasing the risk aversion parameter from 2 to 6 increases the SCC by a mere 5 to 10% even when there are tipping points in the model). I find that puzzling, but it might also simply be true. Requires further research though.

#608 Better representation of the natural cycles of heat, emissions, and ecological processes (e.g., tipping points, methane emissions from thawing permafrost, etc.) Identification of permanent asset damages by climate change

#620 Better damage estimates, especially better treatment of adaptation and extreme weather

#641 less theory, more empirics an end to climate determinism inclusion of the large literature on the links between vulnerability and development inclusion of the large literature on the valuation of health and nature young economists should learn how to use search engine and read the literature, including papers that are more than 5 years old

#644 Invistigate and incorporate individual's perference towards climate change, at present, we only consider the economists'/politicians' perferences.

#650 Quantifying and monetizing more sources of damages (e.g. impacts to marine resources). Ultimately, modeling the climate system with tipping points and feedback is likely to have a larger impact but my sense is that omitted damage sectors are the lower hanging fruit.

#659 Identifying the best studies and focusing mean statistics on these (i.e. removing clearly incomplete or biased studies, as they contaminate the statistics. The topic is too important to leave it to statistics. Informed judgement is important too.

#660 Spend less time dreaming up implausible "tipping point" catastrophes and more time quantifying aerial CO2 fertilization effects, agricultural adaptation strategies and global greening impacts. Also, the authors who look for growth/TFP impacts make glib associations with extreme weather, without citing any data or IPCC assessments on the subject. The links to changes in extreme weather are weak and of ambiguous sign. It doesn't provide a credible rationale for expecting TFP changes due to warming. The use of RCP8.5 and related SSP's should be stopped. It's a ridiculous storyline. Authors who use it are clearly putting their finger on the scale to get a dramatic result and a splashy headline but it is a waste of the reader's time. There's been very little attention paid to Tiebout sorting, but it seems to me it should result in people relocating themselves closer to their privately-optimal temperatures.

#669 Need better process-based models.

#673 Estimates of limited substitutability of non-marked goods Estimates of the persistence of damages on market and non-market sectors with adaptation Robustness analyses with regard to alternative approaches to intergenerational equity and efficiency

#679 Inclusion of additional "damage types" in climate damage functions

#683 Dissagregation. A single number is too abstract to be useful for policy making.

#687 Damages. value of environment/relative prices, One thing that is missing from all of this is the handling of catastrophic risk and risk in general. Issues related to the climate beta, the insurance properties of climate mitigation and so on.

#693 better understanding of climate damages, both the aggregate across time for different climate futures, but also the spatial and socio-economic distribution of damages

#697 Including climate impacts on sentient animals

#708 Better understanding and measuring the pathways through which changes in climate affect the economy (e.g. not just TFP level hits calibrated to reduced form estimates).

#710 ethics, uncertainties-, distribution- and related preferences

#718 Better and more comprehensive modeling of cliamte damages. Better assesment of risks and ambiguities associated with cliamte change.

#738 Damage function

#739 A better estimation of damages, as well as of possibilities and costs of adaptation

#749 I believe we should better try to model the limits of our knowledge

#750 distinguishing between growth and level effects, accounting for non-market damages, damages functions that reflect climate impacts adequately

#752 Getting the valuation of climate (and ecosystems) over an extended future right/or more right in our simplistic economic valuation functions.

Additional feedback

#562 Look forward to the result very much!

#569 This is a really interesting survey!

#570 Bottom line – huge uncertainty over the SCC.

#572 Good luck with your study! This was not an easy survey to complete. I've tried to give you responses that are at least better than just noise :)

#576 This is an interesting survey, though, in some ways, I would have liked a slightly different set up. I wonder if the structural questions should have been upfront to debias individuals. By considering a variety of factors first, you may have better prevented anchoring. If I were to factored in equity weights for example, I think that my response would be very different. In particular, I would have given a range of $50 to $10,000 with a central estimate of $500. Most of the other considerations were factoring into thinking about the appropriate damage function and discount rate, except for substitutability. If I factored this in completely, I would probably have increased by SCC estimates between 50% to 100% even more. In fact, seeing everything laid out in your table about structural assumptions, made me realized how downward biased that I really think that SCC may be despite frequently thinking about the topic.

#592 Good luck with the study. I hesitated to put my name for a moment because i was embarassed if i wrote something contradictory or stupid. But i felt there is a value to being transparent for you. I should say though that answering seriously i should have spent an hour or two. Many of the issues are complex and some questions perhaps could be interpreted in different ways. I am however - as is common every evening, rushed to get through emails and decided to do the questionnaire quickly in 10 minutes and without going back over questions or consulting any document - just straight from the top of my head. I look forward to seing what will come of it and maybe discussing at some point.

#608 It was rather a hard task. I couldn't provide confident numbers myself, but you may still find some patterns from the data when combined with the others' responses.

#620 This is quite possibly the best survey I have ever filled in.

#623 Interesting survey - I'd like to hear the results.

#644 Thanks for your invitation! I think this survey is really useful in terms of gathering ideas.

#687 Tough survey. I am not confident in some of the answers since my reading of the literature is partial and my personal views on the SCC are heavily influenced by my papers and the few papers that I have been influenced by.

#710 It feels VERY wrong to answer this "off the cuff" without thinking longer about it. But some answer is better than no answer (...?). I want to "object" against the framing, at least a little. There is no such thing as an SCC "estimate". We project future marginal damages, at best. And we cannot even do that without several value judgments. To call the SCC an "estimate" may increase its political influence but masks what it is. And makes it vulnerable - you can always critisize any projected value or distribution based on the welfare function parameter choices or structural- or heroic assumptions of any specific model. That's why I really do not like that we use the same terminology as for say a neat RCT estimate of the impact of policy X on behavior Y. Anyway, that was my little pet peeve there. Good luck with the project. PS: Now I'll try to google the actual distribution and bite myself if I am very off ;-)

#738 Good luck with this. Sometimes it was difficult to give a clear cut answer...

#749 One note: I answered your questions as if you were asking for the mean via "central value". Median would be quite different.

#752 I found parts 1 and 3 quite useful. I am afraid that I found part 2 too ambiguous to be meaningfully interpreted.

#1000 I believe the literature has overestimated the SCC. My reasons are the following. Economists are quick to dismiss the benefits of CO2 fertilization. Economists tend to also dismiss future adaptation. Many economists do not realize tipping points are already part of traditional damage estimates. Some economists are reaching for damages that are likely to be small in the labor market, with conflict, from migration, and ecosystem change. Finally, my colleagues are a little confused about the endogeneity of the SCC. If society adopts a high SCC, there will be a lot of mitigation. This will lower the SCC. They should not be calculating an SCC based on Business as Usual unless they believe the SCC will have and should have no effect on policy.

#1001 Thanks for sending me the survey. Honestly, I feel I cannot respond. Sorry for that. Actually, I am now more convinced that the SCC is not helping us and that the cost-benefit analysis of climate change has caused more confusion to the decision-makers that benefit. When Nordhaus obtained the Nobel prize in 2018 the committee included a figure with an "optimal" stabilization temperature of 3C..., something that any IPCC report would endorse. With all my respect, I think the economic methods have been pushed too far in relation to what our data, models and projections and knowledge can really say and not. I have personally worked with the DICE model in the past and I know how sensible it is to small variations to few parameters and on the damage function selected. For that reason, I cannot respond. It can be 10 or it can be 10.000. I don´t know. I am not trying to deviate you from your plans or to argue on this (I might be wrong), but as I have received the mail 4 times, I had the necessity to explain to you why I cannot respond. Of course, I think we need to devote more efforts on putting monetary values to the damage (present and future) of climate change where we have information and keep on with the carbon pricing/markets. Also, if we need a reference-value for guiding decision-makers, I think that the cost of abatement (for example to achieve the Paris Agreement targets) can be a good reference-point as we have less uncertainty on some key technologies. Of course, this is my point of view.

#1002 I looked at your survey, but I can't answer the questions because they are based on the flawed assumption that the current BAU allocation is efficient, which it is not due to the mispricing of Greenhouse Gas Emissions. At an inefficient allocation the Social Cost is not well-defined.

#1003 For building a more complex IAM, I think non-linear climate feedback effects and distributional effects are really important. In the literature on real world impacts of climate change and extremes, it's very clear that the majority of the actual impact comes from climate extremes - heat waves, coastal erosion/inundation and hydrological extremes - which tend to have a non-linear response to climate forcing and there are the huge uncertainties and potentially catastrophic scenarios of runaway feedback mechanisms. In proportional terms, climate change clearly affects

Predictor Variable	Value	Reasoning
SCC Year	Varied	Varied to generate predictions for multiple SCC emission years
All Parametric Variation	1	Allow for any difference in mean SCC due to parametric uncertainty
Backstop	0	True SCC not backstop technology price
Other Market Failure	0	Capture climate externality and not other market failures

Table S10: Random forest prediction variables. Parametric uncertainty is over 14 variables given in Figure 3

the poor a lot more than the rich. That's unlikely to change. I know economists and philosophers go back and forth about the discount rate - what it is and what it should be...but the distributional effects of damages within generations really matter to real people. And if massive inequality of consequence is not adequately addressed, it may lead to an increase in political instability and conflict.

#1007 Thank you for writing. I hope it is ok if I pass on filling in the form. As you probably know, I am not enthusiastic about using IAMs as the basis for calculating social cost of carbon or indeed as a basis for public policy in general. A challenge is to change the structures of economies fundamentally and rapidly. As I have argued, I think we should approach shadow pricing for carbon in a way which helps directly with that question and thus asking what are the prices that, when combined with other structural polices, help make this transformation happen. I fear that the process you are embarking might keep IAMs at centre stage when they really don't capture the urgent and major policy questions.

S.2.3 Random Forest Model and Distribution

The multi-variate analysis described above in Section S.2.1.9 is one way of analysing the variance in literature SCC estimates with a relatively simple, ceteris-paribus interpretation of results. But the method requires selecting a functional form for the regression equation, limiting the extent to which non-linearities or interactions between variables can be examined. Machine-learning models can be used as a complementary analysis to flexibly capture potential non-linearities and interaction effects in settings with many explanatory variables, optimizing for out of sample predictive power. This is particularly valuable in this setting, where the underlying IAMs producing SCC estimates are known to have complex, non-linear relationships between model variables.

We estimate a random forest model on a sample of 1 million draws from our SCC distribution. The distribution used to fit the random forest model samples distributions from the SCC dataset to produce approximately equal representation of all nine major structural model modifications we examine. The model was fit using a dependent variable of log SCC value (requiring us to drop the 2% of observations with a negative SCC value) and 31 predictor variables (9 binary variables representing the presence or absence of a particular structural model characteristic, 14 binary variables representing the presence or parametric variations (distributions) in the model, 2 variables capturing the damage function based SCC estimates, and 6 other variables including discount rate, SCC pulse year etc). The forest is composed of 500 regression trees with a minimum node size of 200 and maximum depth of 12 branches. Feature importance is determined through permutation of individual predictor variables.

For each of the 1 million predictions for each SCC emissions year, we create a vector for predictions from the random forest model based on the SCC year, a draw of the discount rates based on the expert survey in Drupp et al. (*15*), the presence or absence of 9 structural model modifications based on sampling probabilities given by the expert survey in Figure 2 b, and a set of un-varied predictor variables given in Table S10. A random sample of the residuals of the random forest model is added to the predicted value to give a distribution that reflects expert assessments of discount rates and structural model characteristics, variation in the SCC due to structural and parametric uncertainty as reflected in the literature, and residual parametric variation.

Structural Model Characteristic	Probability of Inclusion
Earth System	0.71
Tipping Points: Climate	0.68
Tipping Points: Damages	0.63
Limited Substitutability	0.70
Persistent or Growth Damages	0.62
Epstein-Zin	0.61
Ambiguity or Model Uncertainty	0.60
Learning	0.59
Inequality Aversion	0.63

Table S11: Probability of inclusion in random forest predictions

Probabilities for inclusion or exclusion of a structural model modification in the random forest prediction distribution were based on responses to the expert survey question "To what extent do you agree with the statement that papers including X produce a better SCC than those excluding it?". Probabilities for each structural model characteristic were based on an equal weighting of all survey responses, assigning probabilities of 0,0.25,0.5,0.75 and 1 for the responses "Strongly Disagree", "Disagree", "Neither Agree nor Disagree", "Agree" and "Strongly Agree" respectively, producing the aggregate probabilities give in Table S11.

	SCC Year	Discount Rate
Reference	2045	3.44
	(62)	(1.52)
Earth System	2042	3.21
	(35)	(2.00)
Tipping Points: Climate	2024	3.56
	(27)	(1.32)
Tipping Points: Damages	2023	3.16
	(29)	(1.05)
Persistent / Growth Damages	2027	4.65
	(29)	(1.44)
Epstein-Zin	2034	3.16
	(35)	(1.05)
Ambiguity/Model Uncertainty	2043	4.3
	(36)	(0.94)
Limited Substitutability	2021	2.69
	(20)	(0.78)
Inequality Aversion	2019	4.07
	(27)	(1.85)
Learning	2030	3.08
	(33)	(0.91)

Table S12: Mean and standard deviation (in parentheses) of SCC Year and Discount Rate for datapoints in the reference distribution shown in Figure 2 (estimates without any structural model modifications) and estimates including each of the 10 structural model modifications.

S3 Description of structural model modifications

In this section we describe each of the nine structural model modifications, their relevance to the SCC, and provide example references.

S.3.0.0.1 Ambiguity / Model Uncertainty: (*24, 25*) Climate policy grapples with many kinds and sources of uncertainty, and in many cases there is no agreed-upon quantitative distribution to describe that uncertainty. When multiple possible distributions are available, this problem is referred to as model ambiguity, and solutions typically assume that agents make pessimistic assumptions across the range of possibilities. A standard implementation of model ambiguity would follow a minimax approach, minimizing the welfare loss under the worst set of assumptions.

Higher values of the SCC represent larger losses in welfare for any emissions scenario, so model ambiguity selects higher values of the SCC across any range of possible values. Model ambiguity can also be included when selecting the optimal emission pathway. In this case, model ambiguity results in more cautionary decision-making, lower emissions, and a resulting lower SCC. In this context, the particular role of model ambiguity and the definition of the scenario (exogenous or optimal) is central.

S.3.0.0.2 Earth System (*26, 27*) Studies in this class make changes to the structure of the carbon cycle or warming modules in IAMs compared to the baseline or calibration model. These modules map carbon dioxide emissions into atmospheric concentrations and atmospheric concentrations into climatic changes (principally temperature increases), respectively. Some models collapse the process into a single functional relationship between emissions and temperature.

The Earth system representation matters for the SCC because it determines how much temperatures, and where relevant other climatic variables, respond to CO2 emissions.

S.3.0.0.3 Epstein-Zin Preferences (*28, 29*) Epstein-Zin utility functions disentangle risk across time and states of nature, introducing both a richer set of preferences and the need to solve models recursively. While these types of utility functions have become standard in financial economics over the

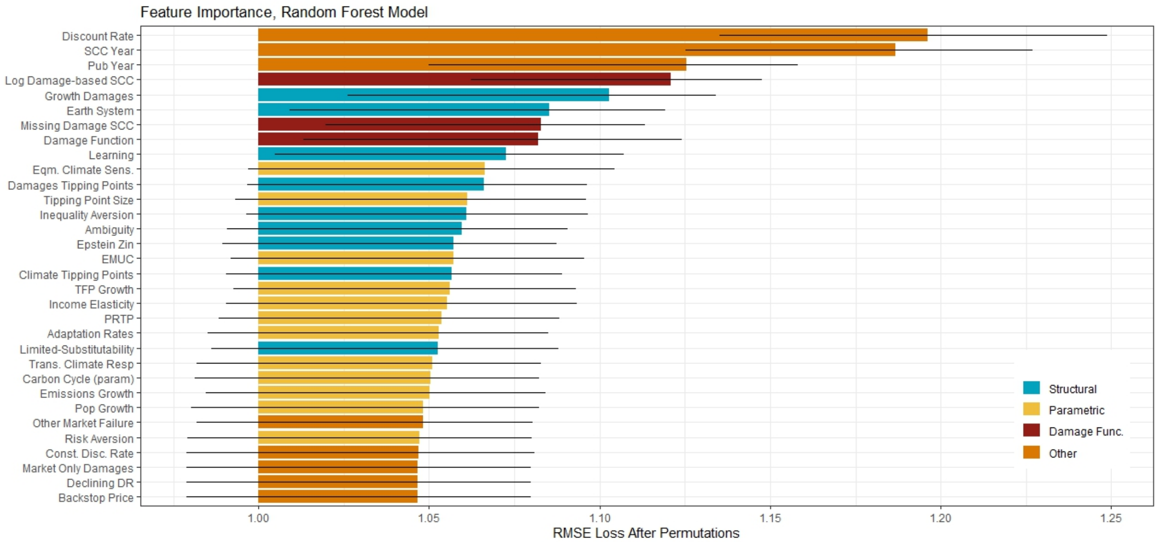

Figure S22: Feature Importance for the random forest model, showing mean change in model testing RMSE after randomly permuting input variables. Error bars show range from permutations.

past 2-3 decades, they are relatively new in climate-economy models. Models that incorporate Epstein-Zin preferences are either computationally involved empirical calibration exercises or highly stylized, attempting to tease out the importance of individual parameters in a set of sensitivity analyses.

Epstein-Zin preferences matter for the SCC, as they ensure that higher risk aversion does not lead to higher discount rates, which in turn lead to lower SCCs. The greater are assumed risks and uncertainties, the greater is the impact of switching to Epstein-Zin preferences on the SCC.

S.3.0.0.4 Inequality Aversion (*30, 31*) The same monetary loss results in a greater welfare loss in poorer regions than it does in richer regions. This intuition is embedded in the concavity of an regional representative agent's utility function or in the region-specific discount factor applied to damages. There is also expected to be a difference between the effect of unequally distributed impacts over space (how welfare losses affect poorer regions) and over time (how welfare losses affect the present, which is assumed to be poorer, from the future).

Populations in hot regions (e.g., near the equator) tend to be poor, and these regions are also expected to have some of the greatest climate impacts. If poorer regions are exposed to greater losses, accounting for inequality aversion can produce large increases in the SCC.

S.3.0.0.5 Learning (*32, 33*) Models that incorporate learning allow the distribution over one or more unknown parameters to evolve over time, typically through a Bayesian update, as data generated by the relationship governed by the unknown parameter are observed over time. Some of the learning models are myopic in that the social planner in the model does not anticipate future learning. Other learning models are forward-looking where the social planner anticipates future learning and adjusts policy accounting for how it affects learning and subsequent welfare. A plurality of the papers in the literature have focused on learning about the equilibrium climate sensitivity – which can be learned from observations of temperature and levels of greenhouse gases – while other papers have included learning about damages, climatic and damage tipping points, and other aspects of the climate economy.

Learning matters for the SCC because as policymakers refine distributions over unknown parameters, they are better able to match climate policy to the true state of the system which will affect the SCC. Learning typically reduces the SCC because of active learning motives where additional carbon emissions magnify the signal of the unknown parameter relative to background noise, allowing for faster learning. Theoretically, learning's effect is a priori ambiguous and depends on nonlinearities and stock effects in the climate-economy.

S.3.0.0.6 Limited Substitutability (*34, 35*) Models that explicitly allow for limited substitutability disaggregate a comprehensive consumption good or a comprehensive capital stock into more detailed component parts. For instance, these model changes feature non-market environmental goods as a direct source of utility and allow for various degrees of substitutability vis-a-vis human-made consumption goods. These models also typically disaggregate the effects of climate damages on both market and non-market goods.

As standard IAMs, like the DICE model, implicitly feature some degree of limited substitutability (Drupp and Hänsel 2021), the effect of explicitly introducing limited substitutability on the SCC is ambiguous. When the degree of substitutability is considered as more limited than implicit in standard models, these model extensions can substantially increase the SCC.

S.3.0.0.7 Persistent / Growth Damages (*4, 36*) Models that allow for persistent or growth damages introduce pathways for transient changes in temperature to have permanent effects. This is usually modeled as temperature affecting the growth rate of factor productivity or increasing the depreciation rate of capital stocks.

Accounting for persistent impacts of temperature generally increases the SCC. If temperature has persistent effects, then a temperature shock today affects output today as well as in all future years. This raises damages and the SCC.

S.3.0.0.8 Tipping Points: Climate (*2, 32*) Climate tipping points/elements have been defined as "subsystems of the Earth system that are at least subcontinental in scale and can be switched – under certain circumstances – into a qualitatively different state by small perturbations" (Lenton et al., 2008). Examples include the permafrost carbon feedback, melting of the Antarctic and Greenland Ice Sheets, and a slowdown of the Atlantic Meridional Overturning Circulation (AMOC). Some models include a representation of the key underlying geophysical relationships that govern the tipping point, while others simulate climate tipping points in a stylized way.

Climate tipping points are diverse in nature and affect the SCC in different ways. Some are positive feedbacks in the climate system that increase the temperature response to CO2 emissions, for example the permafrost carbon feedback. Ice sheet melting increases sea levels and thereby increases coastal adaptation costs and residual costs/damages. Large-scale circulation changes such as AMOC slowdown may primarily change the distribution of impacts across countries and this may affect the overall SCC depending on the incomes of the countries most affected.

S.3.0.0.9 Tipping Points: Damages (*29, 34*) Models that incorporate damage tipping points allow for economic output to be irreversibly reduced, potentially stochastically, if a particular climate threshold is crossed. Damage tipping points are aimed to be stylized ways of capturing events such as melting of the Greenland ice sheet or collapse of the Atlantic Meridional Overturn Circulation. These models incorporate a new state variable that captures the extent of tipping and the probability of this state variable progressing (capturing further advances in tipping and greater damage) depends on the state of the climate system and random shocks.

Damage tipping points matter for the SCC because crossing a tipping threshold irreversibly moves us into a higher-damage world. This tends to increase the SCC because the value of the next ton of carbon must now account for the probability that the tipping threshold is crossed and the subsequent permanent increase in damages. Damage tipping points tend to significantly increase the SCC.

References

1. U. B. of Economic Analysis, Gross Domestic Product: Implicit Price Deflator [GDPDEF] (2021).
2. W. Nordhaus, *Proceedings of the National Academy of Sciences* **116**, 12261 (2019).
3. J. Quiggin, *Australian Journal of Agricultural and Resource Economics* **62**, 4 (2018).
4. K. Ricke, L. Drouet, K. Caldeira, M. Tavoni, *Nature Climate Change* **8**, 895 (2018).
5. E. J. Moyer, M. D. Woolley, N. J. Matteson, M. J. Glotter, D. A. Weisbach, *The Journal of Legal Studies* **43**, 401 (2014).
6. M. L. Weitzman, *Journal of Public Economic Theory* **14**, 221 (2012).
7. M. L. Weitzman, *The Review of Economics and Statistics* **91**, 1 (2009).
8. D. Anthoff, R. S. J. Tol, *Annals of Operations Research* **220**, 223 (2014).
9. K. Gillingham, *et al.*, *Journal of the Association of Environmental and Resource Economists* **5**, 791 (2018).
10. M. Greenstone, E. Kopits, A. Wolverton, *Review of Environmental Economics and Policy* **7**, 23 (2013).
11. D. Anthoff, R. S. J. Tol, Testing the Dismal Theorem (2020).
12. M. N. Conte, D. L. Kelly, *Journal of Environmental Economics and Management* **92**, 677 (2018).
13. S. Ghosh, S. Resnick, *Stochastic Processes and their Applications* **120**, 1492 (2010).
14. M. N. Conte, D. L. Kelly, *Annual Review of Resource Economics* **13** (2021).
15. M. A. Drupp, M. C. Freeman, B. Groom, F. Nesje, *American Economic Journal: Economic Policy* **10**, 109 (2018).
16. M. A. Drupp, F. Nesje, R. C. Schmidt, Pricing carbon, *Tech. Rep. 9608*, CESifo Working Paper (2022).
17. P. H. Howard, D. Sylvan, *Climatic Change* **162**, 213 (2020).
18. R. S. Pindyck, *Journal of Environmental Economics and Management* **94**, 140 (2019).
19. S. DellaVigna, D. Pope, E. Vivalt, *Science* **366**, 428 (2019).
20. J. Hansen, *Environmental Research Letters* **2**, 24002 (2007).
21. R. P. Feynman, Cargo Cult Science (1974).
22. T. Havranek, Z. Irsova, K. Janda, D. Zilberman, *Energy Economics* **51**, 394 (2015).
23. R. S. Tol, *Annu. Rev. Resour. Econ.* **3**, 419 (2011).
24. D. Lemoine, C. P. Traeger, *Journal of Economic Behavior & Organization* **132**, 5 (2016).
25. L. Berger, J. Emmerling, M. Tavoni, *Management Science* **63**, 749 (2017).
26. M. C. Hänsel, *et al.*, *Nature Climate Change* **10**, 781 (2020).
27. S. Dietz, J. Rising, T. Stoerk, G. Wagner, *Proceedings of the National Academy of Sciences* **118** (2021).
28. K. D. Daniel, R. B. Litterman, G. Wagner, *Proceedings of the National Academy of Sciences* **116**, 20886 (2019).
29. Y. Cai, T. S. Lontzek, *Journal of Political Economy* **127**, 2684 (2019).
30. F. Dennig, M. B. Budolfson, M. Fleurbaey, A. Siebert, R. H. Socolow, *Proceedings of the National Academy of Sciences* **112**, 15827 (2015).
31. D. Anthoff, J. Emmerling, *Journal of the Association of Environmental and Resource Economists* **6**, 243 (2019).

32. D. Lemoine, C. Traeger, *American Economic Journal: Economic Policy* **6**, 137 (2014).

33. I. Rudik, *American Economic Journal: Economic Policy* **12**, 340 (2020).

34. Y. Cai, K. L. Judd, T. M. Lenton, T. S. Lontzek, D. Narita, *Proceedings of the National Academy of Sciences* **112**, 4606 (2015).

35. B. A. Bastien-Olvera, F. C. Moore, *Nature Sustainability* **4**, 101 (2021).

36. F. C. Moore, D. B. Diaz, *Nature Climate Change* **5**, 127 (2015).

SCC Meta-Analysis Code Book

Overarching study / value inclusion criteria:

1. Study calculates an “original” SCC (i.e. does not only report values calculated in other studies)
2. Excludes previous meta-analysis estimates
3. Excludes reporting SCC estimates calculated using a model presented in a previous paper. For example, if standard DICE 2007 is run and then modified in some way, the SCC from the modified model should be reported as a new central value, but the value from the standard run should not. (Instead it can be recorded as a “Base Model” value) - this avoids over-counting estimates from models that are re-used by multiple papers
4. Excludes estimates of an “optimal carbon tax” where the primary focus of the paper is modeling market failures other than climate change damages. Optimal carbon tax values that account for market failures other than climate change can be included, but should be flagged in Column M.
5. Excludes estimates of the social cost of methane, N2O or other greenhouse gases / radiatively active species
6. Includes values reported as supplementary analysis in appendix tables

Note that not all values in a paper will necessarily be explicitly coded in the spreadsheet. In some cases, reported variation in the SCC due to parametric variation will be “collapsed” into min-max values and reported as a range in the “SCC Distribution” section. Variation other than that due to SCC year, emissions scenario, discounting scheme, damage function, base / calibration model, or model structure should mostly be reported within a single row.

Bibliographic and Base Model Info

Column A: Unique ID number from full abstract list in “Paper Tracking” spreadsheet

Column B: Add Bibtex reference to this document and record Bibtex name here

Columns C-F: Self-Explanatory

Columns G-H: Report any prior model used for calibration purposes or for comparison. Use Column G (“Base IAM”) if study reports using a prior model and then adding modifications. Use Column H (“IAM Calibrated To”) if study reports calibrating an original IAM using a prior model.

Column I: Categorize the type of SCC study into one of four types:

1. Empirical Improvement: These are studies where the clear goal of the analysis is to improve representation of the climate system, climate damages, or otherwise better align IAM model results with pre-existing scientific or economic findings. Examples might

include improved representation of climate change (e.g. SLR dynamics) or incorporation of econometric damage estimates.
2. Framework Expansion: These are studies where they are adding to or extending the analytical IAM framework, with a view to better representing relevant factors driving climate change costs, but without as strong a tie to a particular empirical literature. Examples could include EZ utility papers, inequality aversion, learning, new discounting frameworks, multi-good utility papers.
3. Sensitivity Analysis: These are studies with a stated intention of only examining the sensitivity of SCC to modeling choices. Often these studies deliberately present a wide range of parameter values and do not take a strong position on which are or are not empirically preferred.
4. Other: Other types of papers can be coded as "Other". Most notably this includes papers reporting "standard" models (e.g. paper reporting the DICE 2013 update and results)

NOTE: This is a categorization of the *paper*, not a particular value from the paper. All rows from the same study should have the same value in this column. If a paper has multiple rows that could be coded differently (i.e. some framework expansion and some empirical improvement) then code the whole paper using the "highest" value, i.e. empirical improvement > framework expansion > sensitivity analysis.

Central SCC Value

Column J: Year of SCC.

- For papers reporting only discrete years, record all SCC years in separate rows.
- For papers reporting continuous SCC values (e.g. graphically), report 2020, 2050 and 2100 if possible. If post-2100 values are reported, also record latest SCC year reported.
- If these specific dates are not available, report earliest available, a mid-century value, and latest available

Column K: Report Central SCC value in $ per ton CO2. If central value is a median, then *also* record it in Column AR (50th percentile).

- Convert values reported in $ per ton C into $ per ton CO2 by dividing by 3.666667
- For values reported in other currencies, convert into $ per ton CO2 using exchange rate from the reported currency year, or the publication year if there is no currency year reported. Report exchange rate in the notes Column (BQ)
- If there is no clear central SCC value (e.g. just high and low values are reported) then this can be left blank and only the range or distribution reported in the "SCC Distribution" section of the sheet.

Column L: Backstop Price? - Enter a 1 if the value recorded in Column K is a backstop price

Column M: Other Market Failure? - Enter a short description of other market failures included in the value in Column K, if that value is an optimal carbon tax that includes market failures other than climate change damages. Examples could include “R&DExternality” or “LimitedRedistribution”.

Column N: Dollar year of SCC - record SCC dollar year if reported.

Column O: Emissions Scenario - record emissions or radiative forcing scenario used for SCC estimate in the row. Examples include:

- Optimal - SCC along the optimal emissions path
- BAU - SCC along the BAU emissions path
- Some radiative forcing scenario (e.g. RCP 7) or socio-economic scenario with associated emissions (e.g. A1B) - assumption is that reported SCC is a BAU SCC along this emissions trajectory (i.e. no mitigation)
- Some temperature or CO2 concentration threshold (e.g. 1.5 degree, 450 ppm) or other implicit emissions constraint (e.g. some risk level). For these cases, enter using the formulation “Constraint - <SPECIFIC CONSTRAINT>” e.g. “Constraint - 2degrees” or “Constraint - 400ppm”. This will help us identify these cases later on in the coding.

-> Only record variation in *emissions* scenario in this column, not variation across other model parameters

Column P: Socio-Economic Scenario - record socio-economic scenario used for SCC estimates in this row. This will particularly be used for determining the consumption growth rate for calculating the Ramsey discount rate. Possible entries include:

- If population and GDP or TFP growth rates are unchanged from the underlying base or calibration model, record the model name here, e.g. “DICE 2007”
- Record standard socio-economic scenarios e.g. SSP2 or A2
- If per-capita consumption growth rate is specified in the paper, record that here (e.g. 2% per year)

Column Q: Reported Base Model SCC - if available, record a Base Model SCC value that is “comparable” to the Central SCC value in the sense of sharing the same 1) discounting assumptions and 2) emissions scenario and 3) SCC year.

Discounting Parameters and Damage Function

Column R: If a constant discount rate (i.e. rather than the Ramsey formula) is used, record the % value here (i.e. 1.5 instead of 0.015 or 1.5%). If a declining discount rate is used, record the initial value here.

Column S: Record Pure Rate of Time Preference. If a declining PRTP is used, record the initial value here.

Column T: If central estimate uses a declining discount rate, add a 1 here. Otherwise leave blank.

Column U: Record EMUC used in Ramsey formula of Central Estimate

Column V-W: Record parameters of Epstein-Zin preferences, if applicable. Note most papers report IES = 1 / EMUC. If necessary convert from EMUC to IES.

Column X: Market Only Damages - Enter a 1 if the Central SCC Value is based on damages that explicitly include only market damages (e.g. calibrated to studies examining variation in GDP)

Column Y: Damage Function Info - Use this column to record the functional form of the damage function for models with only one damage function that depends only on temperature and falls on production (i.e. single region, single sector models). Leave blank if the model is not of this type (e.g. multi-region models). Information can be recorded in one of three ways:

1. If the damage function is left unchanged from the base or calibration model, re-enter that model name in here (eg. DICE 2007)
2. If the damage function is a variation commonly referred to by the original authors, enter the corresponding name here. Possible entries are:
    a. Weitzman - corresponding to damage function used in Weitzman (2012, Journal of Public Economics Theory)
    b. HowardSterner - corresponding to base damage function reported in Howard and Sterner (2017, Environmental and Resource Economics). (Note - some HowardSterner damage function specifications also include calibration to growth rate impacts, which can be indicated using "Calibrated" in Column AD)
    c. DietzStern - corresponding to damage function in Dietz and Stern (2015, The Economic Journal)
3. If neither 1 nor 2 apply, directly enter the damage function into the column, using T as the GMST change from the baseline used in the paper (e.g. 0.0004 * T^3). Any R expression is permitted. T is assumed to be contemporaneous temperature (i.e. warming in the same time period as damages occur) unless otherwise indicated, using syntax T_{t-k}.

Structural Changes

Column Z: Carbon Cycle - Enter a 1 if the values in this row include a structural change in the Carbon Cycle model, compared to the baseline or calibration model, or compared to DICE if no baseline or calibration model is reported.

Column AA: Climate Model - Enter a 1 if the values in this row include a structural change in the Climate Model (i.e. effects on the physical climate system, conditional on greenhouse gas concentration, including temperature and sea-level rise), compared to the baseline or calibration model, or compared to DICE if no baseline or calibration model is reported.

Column AB: Climate Tipping Points - Enter a 1 if the values in this row include climate tipping points. This means a representation of specific changes in the earth system such as ice-sheet processes, Amazon dieback, changes to the thermo-haline circulation.

Column AC: Damage Tipping Points - Enter a 1 if the values in this row include a stylized or abstract representation of tipping points as a change in damages without modeling the underlying drivers from the climate system (e.g. Cai, Lonztek JPE).

Column AD: Persistent / Growth Damages - Enter a 1 here if the values in this row include persistent damages, for instance via damages to the capital depreciation rate, TFP growth, or capital stock. Enter “Calibrated” if the damage function is calibrated to partially account for persistent damages but they are not represented structurally in the model.

- Note that standard DICE, because of endogenous capital formation, includes some small persistence in damages. Do NOT enter a 1 here for standard versions of DICE.

Column AE: Epstein Zin Utility - Enter a 1 if the values in this row come from a model using Epstein-Zin preferences

Column AF: Ambiguity / Model Uncertainty - Enter a 1 if the values in this row explicitly account for ambiguity or model uncertainty

Column AG: Non-Substitutable Goods - Enter a 1 if the values in this row come from a model with more than one good in the utility function that are imperfectly substitutable with each other

Column AH: Inequality Aversion - Enter a 1 if the values in this row include inequality aversion. Enter “Calibrated” if the damage function is calibrated to account for inequality aversion but it is not represented structurally in the model.

Column AI: Learning - Enter a 1 if the values in this row are from a model explicitly representing a learning process.

Column AJ: Alternative Ethical Approaches - Enter a 1 if the values in this row are based on a different ethical approach than discounted utilitarianism.

SCC Distribution

Columns AK to BA

Use these columns to record any reported variation around the Central Value recorded in Column K of this row. Values recorded in these columns should share the same SCC year, emissions scenario, and structural changes as the recorded central value. Unless damage function parameters and / or discounting parameters are being varied, they should also share the same damage function and discount rate.

Use “Min” (Column AK) and “Max” (Column BA) to report ranges with the two most extreme SCC values at the upper and lower end resulting from parameter changes not associated with probabilities

Use quantile columns to enter any reported quantiles of the distribution. This could come either from quantiles reported from a Monte Carlo sampling, from reported confidence intervals around values (e.t. 95% confidence intervals), or from one at a time sensitivity analysis, where the varying parameters have quantiles associated with them (e.g. varying the climate sensitivity to +1 or -1 standard deviation).

Note that the recorded distribution must be “well behaved” - i.e. values must be strictly increasing with quantiles. Otherwise the distribution sampling algorithm will not be able to interpret the entries.

All units in $ per ton CO2. If necessary, convert from $ per ton C or from values reported in other currencies, consistent with recording central value.

If central value is a median, *also* record value in the 50th percentile column (column AS)

Parametric Uncertainty

Use these columns to record parametric variation giving rise to either:

1) Variation recorded in the SCC distribution block. Enter a 1 if variation in this parameter contributes to the range reported in Columns AJ to BA

   OR

2) If no distribution is reported, but the central value recorded in Column K is a central value from a distribution (without the distribution being reported), then add 1s in the relevant columns in this section. For example, most PAGE09 values come from Monte Carlo runs of PAGE that include variation in:
   a) Transient climate response
   b) Carbon cycle
   c) Tipping Point Magnitude
   d) Damage Function
   e) Adaptation Rates
   f) Damage Income Elasticity

Additional Information

Column BP: Paper Location - Record where in the paper values in the row come from (e.g page number, figure, table). Could be multiple locations if central value and parametric variation ranges come from different places

Column BQ: Flag - Use this to add other informational flags to aid interpretation of values in this row

Column BR: Notes - Add any other notes here aiding interpretation of the values in this row

- Any additional notes on decisions made in coding a paper or information required to interpret values can be added in this Google Doc.