ECONOMICS

Working Paper Series

No. 19-2020

Testing the Dismal Theorem

David Anthoff
Energy and Resources Group, University of California

Richard S. J. Tol
Department of Economics, University of Sussex
r.tol@sussex.ac.uk

Abstract: Weitzman's Dismal Theorem has that the expected net present value of a stock problem with a stochastic growth rate with unknown variance is unbounded. Cost-benefit analysis can therefore not be applied to greenhouse gas emission control. We use the Generalized Central Limit Theorem to show that the Dismal Theorem can be tested, in a finite sample, by estimating the tail index. We apply this test to social cost of carbon estimates from three commonly used integrated assessment models, and to previously published estimates. Two of the three models do not support the Dismal Theorem, but the third one does for low discount rates. The meta-analysis cannot reject the Dismal Theorem.

JEL classification: C46, D81, Q54

Key words: climate policy; dismal theorem; fat tails; social cost of carbon

Testing the Dismal Theorem

David Anthoff[a], Richard S.J. Tol[b,c,d,e,f,g,1,*]

[a] *Energy and Resources Group, University of California, Berkeley, USA*

[b] *Department of Economics, University of Sussex, Falmer, United Kingdom*

[c] *Institute for Environmental Studies, Vrije Universiteit, Amsterdam, The Netherlands*

[d] *Department of Spatial Economics, Vrije Universiteit, Amsterdam, The Netherlands*

[e] *Tinbergen Institute, Amsterdam, The Netherlands*

[f] *CESifo, Munich, Germany*

[g] *Payne Institute for Earth Resources, Colorado School of 36 Mines, Golden, CO, USA*

Abstract

Weitzman's Dismal Theorem has that the expected net present value of a stock problem with a stochastic growth rate with unknown variance is unbounded. Cost-benefit analysis can therefore not be applied to greenhouse gas emission control. We use the Generalized Central Limit Theorem to show that the Dismal Theorem can be tested, in a finite sample, by estimating the tail index. We apply this test to social cost of carbon estimates from three commonly used integrated assessment models, and to previously published estimates. Two of the three models do not support the Dismal Theorem, but the third one does for low discount rates. The meta-analysis cannot reject the Dismal Theorem. *Keywords*: climate policy; dismal theorem; fat tails; social cost of carbon
JEL codes: C46, D81, Q54

1. Introduction

The Dismal Theorem (Weitzman, 2009a) has that the uncertainty about climate change is too large for expected utility maximisation. Specifically, Weitzman showed that the expected value of the social cost of carbon, the marginal net present impact of greenhouse

* Jubilee Building, BN1 9SL, UK
Email address: r.tol@sussex.ac.uk (Richard S.J. Tol)
URL: http://www.ae-info.org/ae/Member/Tol_Richard (Richard S.J. Tol)

[1] We fondly remember lengthy conversations with Marty Weitzman on this topic. Tommi Ekholm, Samuel Okullo, Antonin Pottier, and Gernot Wagner graciously shared their estimates of the social cost of carbon. Bill Nordhaus had good suggestions on the tail-index estimator.

gas emissions or, if evaluated along the optimal emissions trajectory, the Pigou (1920) tax, is unbounded. Weitzman's is an analytical result for a highly stylized model.[2] We here test for the existence of the first moment of the social cost of carbon in three integrated assessment models and in a meta-analysis of previously published estimates. We find that, under more realistic assumptions, the Dismal Theorem stands.

We test the Dismal Theorem by estimating the tail-index of social cost of carbon distributions. The Generalized Central Limit Theorem holds that the distribution of a sum (and so the mean) of independent, identically distributed variables tends to a stable distribution, and that its right tail will follow a power law with index α. If the variance of the summands is finite, $\alpha = 2$ and the distribution of the sum is Gaussian—this is the Central Limit Theorem, ungeneralized. $\alpha < 2$ indicates that the variance is infinite, $\alpha < 1$ that the average does not exist. Weitzman's Dismal Theorem can thus be restated as: The tail-index of the social cost of carbon is less than one.

Weitzman (2009a) pointed out that the expectation of the discount factor is the Moment Generating Function (MGF) of the discount rate. If the growth rate of the economy is Normally distributed with unknown variance, then the estimated growth rate has a Student t distribution. In a stroke of genius, Weitzman combined these two basic insights to show that the MGF of the social cost of carbon does not exist.

The MGF is an alternative specification of a Probability Distribution. It can be thought of as an expansion in its moments: The n$^{\text{th}}$ partial derivative of the MGF equals the n$^{\text{th}}$ moment. The MGF thus specifies mean, variance, skewness, kurtosis and all higher moments, and reveals whether or not these moments exist. The tail-index similarly reveals which moments exists. Tail-index and MGF are thus connected. However, while we can estimate the MGF from the empirical MGF $M(t) = n^{-1} \sum_j e^{tx_j}$, in a finite sample, empirical moments are by definition finite. We would then need to design a test whether the empirical moments are truly finite.[3] On the other hand, there are ready-made estimators of the tail-index, also for small α. Estimating the tail-index thus provides a direct test of the Dismal Theorem. We refer to tails with $\alpha < 1$ as *fat*, we call tails with $1 \leq \alpha < 2$ *thick*, while *thin* tails have $\alpha \geq 2$.

We estimate the tail-index of the social cost of carbon for three Integrated Assessment Models—DICE, FUND, and PAGE—that have been previously used to estimate the social cost of carbon, including by the Obama administration (on the Social Cost of Carbon, 2010, 2013). These models are implemented in MIMI, which allows for parametric uncertainty analysis, scenario analysis, and structural uncertainty analysis.

We further estimate the tail-index for the population of published estimates of the social cost of carbon. These estimates use the above integrated assessment models but also other

[2]See Cato (2020) for an even simpler version that can be used in class: The Dismal Theorem is the dual of the St Petersburg Paradox.

[3]Anthoff and Tol (2014) repurpose stationary tests for this.

models, thus extending the analysis of the impact of model structure on the social cost of carbon and the tail of its distribution. The published estimates were screened by authors and referees prior to publication, reducing the risk that estimates in the tail of the distribution are the result of the mechanical extrapolation that sometimes happens in Monte Carlo exercises.

This is not the first paper to test the Dismal Theorem. Millner (2013) notes that the Dismal Theorem holds if relative risk aversion is constant but not if absolute risk aversion is hyperbolic, echoing Geweke (2001). Millner (2013) further finds that the Dismal Theorem does not stand if greenhouse gas emission reduction is added to the model. Horowitz and Lange (2014) show that the Dismal Theorem holds in partial but not in general equilibrium: Changes in savings and investment would also prevent a collapse of the economy.

Anthoff and Tol (2014) take the Dismal Theorem literally—you cannot apply expected cost-benefit analysis to climate policy—and explore alternative decision criteria. They also design a statistical test for the fatness of the tail. This test relies on the recursive estimates of net present welfare for an expanding Monte Carlo sample. They show that, in FUND, mean net present welfare does not converge, violating the Law of Large Numbers and thus the Central Limit Theorem—as predicted by the Dismal Theorem. However, their statistical test—Augmented Dickey-Fuller—was designed for a different purpose. In this paper, we use statistical tests explicitly designed for the tail-index. We also use a wider range of models.

The paper proceeds as follows. We discusses the models, data and statistical tests used in Section 2. More details are presented in the Appendix. Section 3 presents and discusses the results. Section 4 concludes.

2. Methods

2.1. Integrated Assessment Models

Integrated Assessment Models (IAMs) come in many shapes and forms. The ones used in this paper combine representations of population, economic output, energy use, greenhouse gas emissions, carbon cycle, climate, and impacts of climate change. These models were designed to inform the optimal course of action on greenhouse gas emission reduction either by maximising net present welfare or, approximately, equating the marginal costs and benefits of emission abatement. The three models used in this paper were selected to advice the Obama Administration on the social cost of carbon, the internal carbon price used for regulatory cost-benefit analysis.[4]

2.1.1. DICE

DICE is the oldest and most prominent of IAMs (Nordhaus, 1992, 1993, 2018). It extends a Ramsey/Cass/Koopmans model of economic growth with a Maier-Reimer/Hasselmann model of the carbon cycle and a Schneider/Thompson model of climate change. The impacts

[4]The Trump Administration uses the same three models, but limits the social cost of carbon to impacts that fall on the USA and uses a higher discount rate than the previous administration.

of climate change are given by a power function. Energy is modelled as a derived demand. Emission reduction costs are another power function. The model solves optimal investment and optimal emission abatement.

2.1.2. FUND

FUND uses scenarios of economic growth rather than a growth model. The carbon cycle and climate parts are very similar to DICE (Tol, 1999). Instead of a single damage function, FUND models impacts of climate change separately for each sector (Tol, 2002a,b). The model can be used to solve optimal emission abatement, but also to estimate impacts along arbitrary emission or climate scenarios (Anthoff et al., 2016).

2.1.3. PAGE

PAGE has an economy as simple as FUND and climate change impacts as simple as DICE. Carbon cycle and climate are similar to those in the other two models (Plambeck et al., 1997, Hope and Schaefer, 2016). The key strength of the PAGE model is that it is centred on the analysis of parametric uncertainty.

2.1.4. MIMI

Mimi is a Julia package that splits large models like DICE, FUND and PAGE into smaller components, as recommended by National Academies of Sciences and Medicine (2017). Each of the component corresponds to a well-delineated part of the cause-effect chain, often with strong disciplinary roots—the carbon cycle, for instance, or the impacts of climate change on heat stress. Each component can be tested and used separately, and combined in any logical permutation with other components. See Moore et al. (2018). For instance, MIMI makes it easy to combine PAGE's climate model with DICE's impact function, or replace FUND's carbon cycle without reprogramming the entire model. MIMI also offers facilities for Monte Carlo analysis and data management. It is those features we use in this paper.

2.2. Meta-analysis

There is a large literature on the social cost of carbon spanning four decades, from Nordhaus (1982) to Okullo (2020). Tol (2020) updates a previous meta-analysis of the social cost of carbon (Tol, 2018). He counts 2786 estimates in 148 papers. These are estimates of the social cost of carbon of carbon dioxide emitted in the recent past.

The estimates are treated in three different ways. First, all estimates are treated equally. Second, estimates are weighted such that the total weight *per paper* equals one. Within each paper, estimates that are favoured by the authors are given higher weight than estimates that are presented for robustness or replication. Third, papers are weighted by quality, as measured by peer-review, age, scenario use, completeness, and mathematical consistency.

2.3. The Tail-Index and its Estimation

Hill (1975) derives the Maximum Likelihood estimator for the tail-index and shows that it is unbiased. If the right tail of distribution follows a power law, then its natural logarithm is a straight line. Hill estimates the slope of that line. A key consideration is the definition of the tail. We here follow the literature and estimate the tail-index for the 10 largest observations up to the 10% largest observations. In addition, we use the Huisman et al. (2001) estimator that combines estimatess across tail-sample sizes.

Much of the literature that follows Hill (1975) is concerned with estimators in cases where the tail only approximately follows a power law. Many estimators have been proposed (Fedotenkov, 2018). We here use those with properties that are widely accepted: maximum likelihood (ML), best linear unbiased (BLUE), least squares (LS), method of moments (MM), and quantile-quantile (QQ). Furthermore, we estimate, using numerical maximum likelihood, the tail-index for a specific deviation from the Pareto Distribution: The Generalized Pareto Distribution. Details are given in Appendix A.

3. Results

3.1. Integrated Assessment Models

Estimates of the social cost of carbon are based on a large number of parameters, all of which are uncertain or disputed, if not very uncertain or controversial. Parametric uncertainty is here reflected by Monte Carlo analysis, which of course replaces an assumption about the value of a particular parameter with assumptions about its distribution and moments. For the sake of space, we do not test every assumption in a sensitivity analysis. Instead, we select a core set of assumptions, and vary the key parameters one or two at a time. The pure rate of time preference is set to 1% per year, with 0.1% and 3% as sensitivities. The rate of risk aversion is set to 1, with 0.5 and 1.5 as alternatives. Discounting follows the Ramsey Rule, but as the projected growth rate is uncertain, the certainty-equivalent discount rate falls with the time horizon. As an alternative, we use a constant consumption discount rate, as was done in the official US federal social cost of carbon estimates. In addition, we also test social cost of carbon estimates that use equity weighting. In those estimates we set the inequality aversion parameter equal to the chosen risk aversion parameter. In graphs, we show all six estimators of the tail-index for a range of tail-sample sizes. In tables, we show the Huisman-type estimators that synthesize across tail-sample sizes.

All these choices are debatable, but making a choice is preferred to showing results for every permutation of parameters, models, and estimators. Data and code are available to the reader who wants to test alternative choices.

Weitzman (2009b, see also Weitzman (2010)) notes the key distinction between *additive* and *multiplicative* damages. If U denotes utility, C consumption, and D damage, for multiplicative damages, $U = U(C(1-D)^{-1})$ so that $U \to -\infty$ for $D \to \infty$. For additive damages, $U = U(C - D)$ so that $U \to -\infty$ for $D \to C$. In either representation, there is an upper limit to damage, $D \leq D_{\max}$, but the maximum willingness to pay to avoid negative impact is everything you earn plus everything you own and can borrow, $D_{\max} > C$. Additive damages are thus much more likely to lead to large utility losses than multiplicative damages. DICE has multiplicative damages, FUND and PAGE have additive damages.

Table C.7 shows the mean and standard deviation of the social cost of carbon for the 100,000 runs in the Monte Carlo analysis with the three integrated assessment models, for a pure rate of time preference of 1% and a rate of risk aversion of 1. PAGE is the most pessimistic and FUND the most optimistic, with DICE in between. However, DICE is the most confident and PAGE the least, which FUND in between. Mean and standard deviation do not fully describe the characteristics of the Monte Carlo results. Figures C.6, C.7 and C.8 show the probability densities for the whole sample and the 1000 largest observations for the three models. The probability densities are unimodal and right-skewed, but the similarity between models ends there. The density for DICE looks like a lognormal distribution, while PAGE is much like an exponential distribution. FUND has a more pronounced left tail—this model explicitly accounts for savings in winter heat costs, for avoided cold-related deaths, and for carbon dioxide fertilization—and a much thicker right tail than the other two models.

3.1.1. DICE

Figure C.6 shows the histogram of all 100,000 Monte Carlo runs, and for the 1,000 largest. Figure 1 shows results for the tail-index for DICE. The rate of risk aversion is 1, the pure rate of time preference 1%. All six estimators are shown.

The moments estimator is unreliable, moving from very high to very low estimates and back. The Generalized Pareto puts the tail-index around zero, but the maximum likelihood estimators warns that these results are unreliable. These are signs of a thin tail, as shown in Appendix A.

The other estimators are more robust to deviations from the null hypothesis. These estimators indicate that the social cost of carbon has a thin tail: The tail-index is greater than two, both mean and variance exist. In fact, the first seven moments exist. The quantile-quantile estimator shows a thinner tail than the three analytic estimators. The tail thins somewhat if we consider fewer, more extreme observations.

Table 1 gives Huisman-type estimates for a range of parameters. Apart from the unreliable moments and Generalized Pareto estimators, estimates of the tail-index are significantly greater than two, regardless of the pure rate of time preference and the regardless of the rate of risk aversion.

3.1.2. FUND

Figure 2 shows the six estimators of the tail-index for a range of tail-sample sizes for FUND, with risk aversion equal to 1 and the pure rate of time preference equal to 1%, using equity weights.

In the deep tail, the tail-index is smaller than one and the Dismal Theorem holds. If we consider more than 0.5% of the largest observations, the tail is thick but not fat. According to the maximum likelihood, least squares and best linear unbiased estimators, the tail thins quickly as we add less extreme observations. The moments and quantile-quantile estimator show less rapid thinning. The Generalized Pareto estimator shows very rapid thinning followed by a thickening.

	ML	BLUE	LS	MM	QQ	GP
0.5 0.1%	7.8856	7.8366	7.7385	-133.0361	8.8574	-0.14203
	(1.2766e-21)	(1.2504e-21)	(1.1864e-21)	(3918.7396)	(3.1179e-23)	(0.067771)
1.0 0.1%	6.5847	6.5464	6.4697	11.6932	6.6153	0.086826
	(1.7019e-18)	(1.6862e-18)	(1.707e-18)	(8.5067)	(4.4255e-20)	(2.6108)
1.5 0.1%	7.1006	7.0593	6.9768	39.6014	7.2656	0.063021
	(1.7728e-18)	(1.7987e-18)	(1.8672e-18)	(39.5804)	(1.5567e-20)	(0.053028)
0.5 1.0%	10.2106	10.1519	10.0343	207.6101	10.6491	-0.011441
	(1.605e-25)	(1.6559e-25)	(1.7891e-25)	(1519.12)	(4.3116e-29)	(0.046044)
1.0 1.0%	8.9743	8.9197	8.8106	-1752.9251	9.6979	-0.078542
	(9.2298e-26)	(9.015e-26)	(8.5354e-26)	(68527.6968)	(2.1026e-27)	(0.054023)
1.5 1.0%	8.957	8.9042	8.7987	-1172.9642	9.3903	-0.0053218
	(6.6621e-25)	(6.8934e-25)	(7.4579e-25)	(16141.3101)	(6.8281e-27)	(0.06128)
0.5 3.0%	15.9737	15.8755	15.6791	-0.18268	17.7286	-0.078184
	(3.849e-41)	(4.0854e-41)	(4.6035e-41)	(18.6944)	(8.5954e-47)	(0.036076)
1.0 3.0%	17.1236	17.0217	16.818	1488.2843	17.7648	-0.04104
	(1.7807e-46)	(1.8744e-46)	(2.0844e-46)	(37134.2168)	(6.9672e-52)	(0.04387)
1.5 3.0%	18.5796	18.4715	18.2554	166.6744	19.2991	-0.036273
	(1.2876e-48)	(1.3827e-48)	(1.6153e-48)	(1995.5483)	(5.1082e-55)	(0.027032)

Table 1: Huisman-type estimates of the tail-index, for different rates of risk aversion and pure time preference (rows), and for different estimators of the tail-index (columns). All results use DICE. Equity weights are inapplicable in a one-region model.

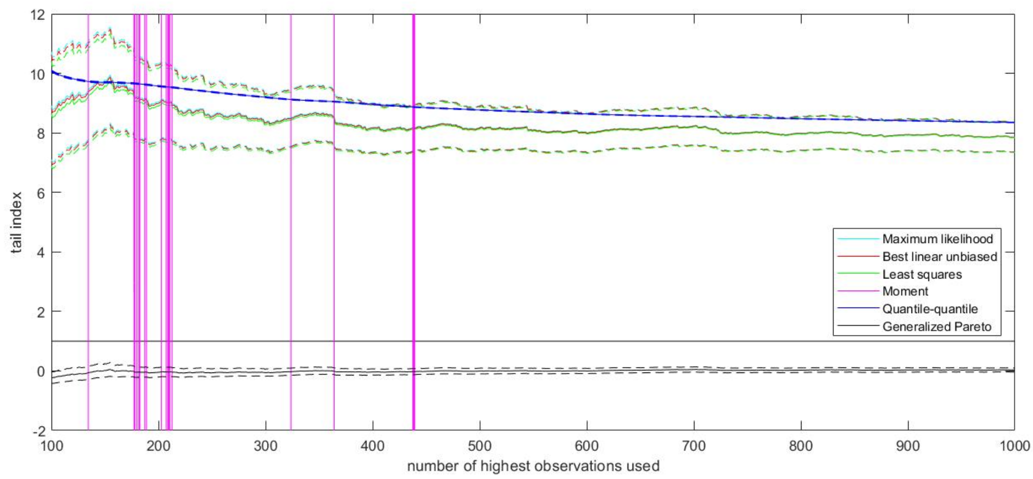

Figure 1: Estimates of the tail-index for DICE, for a pure rate of time preference of 1% per year, and a rate of risk aversion of 1.

	ML	BLUE	LS	MM	QQ	GP
0.5 0.1%	0.49929	0.49465	0.48536	0.75055	0.73263	1.2793
	(0.00043394)	(0.00043034)	(0.00042331)	(0.00046648)	(0.00045236)	(2.4503e-05)
1.0 0.1%	0.54059	0.53531	0.52476	0.79267	0.81602	1.3192
	(9.737e-05)	(9.6414e-05)	(9.4486e-05)	(3.3557)	(0.00026509)	(0.40182)
1.5 0.1%	0.66252	0.65631	0.64388	1.0628	0.9545	1.0139
	(3.2906e-05)	(3.2304e-05)	(3.1063e-05)	(0.00014926)	(9.5983e-05)	(0.0011122)
0.5 1.0%	0.4902	0.48516	0.47509	0.78699	0.72682	1.6117
	(4.5477e-05)	(4.4635e-05)	(4.298e-05)	(0.00031624)	(0.00029613)	(0.00035186)
1.0 1.0%	0.62409	0.61793	0.60561	0.90613	0.78175	1.7721
	(1.6716e-05)	(1.7677e-05)	(1.988e-05)	(6.4247e-05)	(8.1921e-05)	(0.00057725)
1.5 1.0%	1.3976	1.3873	1.3668	1.3408	1.1974	1.3021
	(6.068e-07)	(6.7601e-07)	(8.4227e-07)	(4.3422e-06)	(5.7908e-07)	(0.0097273)
0.5 3.0%	2.0845	2.0716	2.046	1.2473	1.0355	1.4032
	(7.1847e-07)	(9.5222e-07)	(1.8772e-06)	(3.8228e-06)	(4.2536e-06)	(0.025589)
1.0 3.0%	3.4624	3.4442	3.4079	1.4224	1.4062	1.1857
	(2.9351e-07)	(4.2401e-07)	(9.3968e-07)	(5.5725e-07)	(5.0027e-07)	(0.16091)
1.5 3.0%	6.6016	6.5649	6.4913	3.7086	5.8254	0.2220
	(5.1705e-19)	(5.7713e-19)	(7.2668e-19)	(1.0345e-16)	(1.2135e-19)	(2.1105)

Table 2: Huisman-type estimates of the tail-index, for different rates of risk aversion and pure time preference (rows), and for different estimators of the tail-index (columns). All results use equity-weights. All results use FUND.

Table 2 shows the results for the same model for a range of rates of risk aversion and pure rates of time preference. The tail thins as the future is discounted harder. This is as expected: The greatest uncertainties lie in the further future, not just because less is known about the more distant future, but also because climate change is more pronounced later.

3.1.3. PAGE

Figure 3 shows results for PAGE. The rate of risk aversion is 1, the pure rate of time preference 1%. Equity weights are used. All six estimators are shown.

The moments estimator is unreliable, jumping between very high and very low estimates. The Generalized Pareto puts the tail-index around zero, but it too is unreliable judging from the warning messages from the numerical optimization algorithm. These are signs that the tail is thin; see Appendix A.

The other estimators are less fragile to deviations from the null hypothesis that the tail follows a power law. These estimators show that the social cost of carbon has a thin tail: The tail-index is greater than two, both mean and variance exist. The quantile-quantile estimator shows a thinner tail than the three analytic estimators. The deeper tail is thinner, as PAGE uses triangular distributions, capped from below and above, to describe parametric uncertainties.

Table 3 gives Huisman-type estimates for a range of parameters. Apart from the unreliable

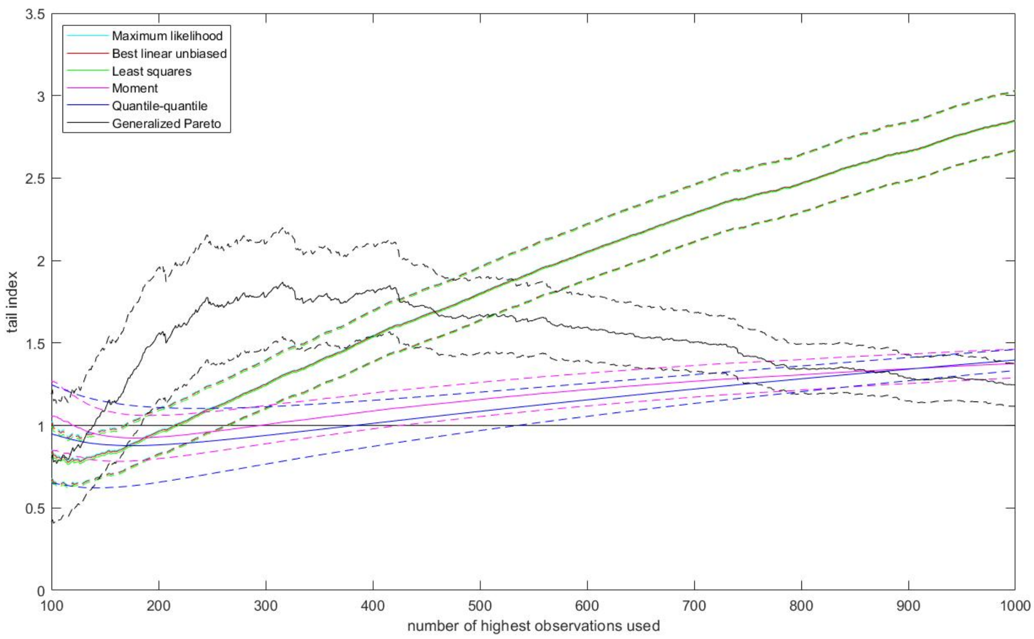

Figure 2: Estimates of the tail-index for FUND, with equity weights, a pure rate of time preference of 1% per year, and a rate of risk aversion of 1.

moments and Generalized Pareto estimators, estimates of the tail-index are significantly greater than two, regardless of the pure rate of time preference and the regardless of the rate of risk aversion.

3.1.4. Equity weights

Figure 4 repeats Figure 2 but without equity weights. Above, we use Pearce equity weights (Fankhauser et al., 1997, see Anthoff et al. (2009a) and Anthoff and Tol (2010) for alternatives), weighing impacts by the ratio of global to regional income raised to the power of the rate of risk aversion. The pattern without equity weights is similar to the pattern with, but the tail-index is somewhat larger, that is, the tail of the social cost of carbon is somewhat thinner. Equity weights emphasize the impact on the poor, who tend to be more vulnerable. Equity weights thus amplify the risks of climate change. Numerically, however, the effect on the tail-index is small.

Table 4, which repeats Table 2, confirms that, quantitatively, equity-weights have a small effect on the estimated tail-index. However, equity-weights can fatten as well as thin the tail—compare the bottom rows of Tables 2 and 4. The tail thins for larger discount rates because, in the short-run more carbon dioxide in the atmosphere stimulates agriculture through carbon dioxide fertilization, which primarily benefits poorer countries (Anthoff et al., 2009b).

	ML	BLUE	LS	MM	QQ	GP
0.5 0.1%	3.5171	3.4971	3.4572	11.1351	3.8611	0.095837
	(5.7688e-06)	(5.6035e-06)	(5.3235e-06)	(41.8374)	(3.6028e-08)	(0.057553)
1.0 0.1%	3.578	3.5568	3.5145	18.4657	4.1452	0.0018793
	(7.1365e-06)	(6.9361e-06)	(6.5526e-06)	(105.1644)	(5.9296e-08)	(0.038898)
1.5 0.1%	3.5539	3.5328	3.4906	-1.3102	4.0503	0.025607
	(5.255e-06)	(5.0219e-06)	(4.5636e-06)	(264.198)	(4.0186e-08)	(0.026274)
0.5 1.0%	3.224	3.205	3.1671	10.9081	3.6503	0.065545
	(5.7691e-06)	(5.6211e-06)	(5.3445e-06)	(2.9763)	(8.3598e-08)	(0.034658)
1.0 1.0%	3.4873	3.4657	3.4224	71.4206	4.1751	-0.10081
	(3.9395e-06)	(3.8006e-06)	(3.5275e-06)	(1007.0128)	(6.5447e-08)	(0.055857)
1.5 1.0%	3.3435	3.3243	3.2859	9.1511	3.6876	0.094431
	(6.081e-06)	(5.8576e-06)	(5.4579e-06)	(9.5772)	(5.2809e-08)	(0.041587)
0.5 3.0%	3.4146	3.3954	3.3568	6.3418	3.6801	0.12975
	(5.4577e-06)	(5.2277e-06)	(4.8161e-06)	(2.523)	(3.2468e-08)	(0.038091)
1.0 3.0%	3.4409	3.4215	3.3827	9.4061	3.7268	0.1390
	(7.9737e-06)	(7.8294e-06)	(7.6052e-06)	(5.4386)	(3.2242e-08)	(0.049278)
1.5 3.0%	3.367	3.3455	3.3024	6.592	4.0417	-0.10817
	(6.963e-06)	(6.618e-06)	(5.9228e-06)	(208.6886)	(1.4336e-07)	(0.066987)

Table 3: Huisman-type estimates of the tail-index, for different rates of risk aversion and pure time preference (rows), and for different estimators of the tail-index (columns). All results use equity-weights. All results use PAGE.

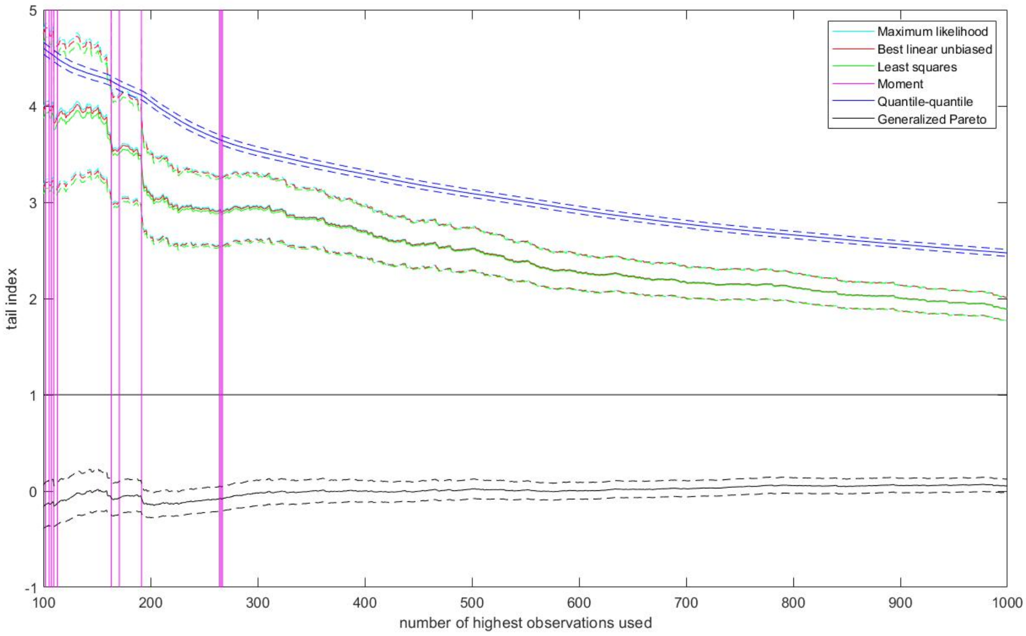

Figure 3: Estimates of the tail-index for PAGE, with equity weights, a pure rate of time preference of 1% per year, and a rate of risk aversion of 1.

	ML	BLUE	LS	MM	QQ	GP
0.5 0.1%	0.57278	0.56774	0.55765	0.84208	0.80036	1.1252
	(0.00022036)	(0.00021603)	(0.00020736)	(0.00047405)	(0.00036518)	(0.20515)
1.0 0.1%	0.61283	0.6073	0.59624	0.86077	0.75319	1.6733
	(1.9223e-05)	(1.999e-05)	(2.1835e-05)	(0.61221)	(0.0002429)	(0.31735)
1.5 0.1%	0.7972	0.79087	0.77821	0.9481	0.85251	1.3086
	(2.0684e-06)	(2.1651e-06)	(2.456e-06)	(1.7998e-05)	(9.87e-06)	(0.00052268)
0.5 1.0%	0.57662	0.57094	0.55957	0.90451	0.81242	1.5341
	(1.0601e-05)	(1.0526e-05)	(1.0432e-05)	(0.13496)	(0.00013606)	(0.37771)
1.0 1.0%	1.0023	0.99439	0.97851	1.0706	0.89511	1.7289
	(6.8892e-06)	(7.615e-06)	(9.3622e-06)	(0.15092)	(8.5887e-06)	(0.25685)
1.5 1.0%	1.2932	1.2841	1.2659	1.0561	0.83036	1.6055
	(5.5448e-06)	(6.6525e-06)	(1.0402e-05)	(1.203e-05)	(1.1392e-05)	(0.0034558)
0.5 3.0%	2.3585	2.3454	2.3192	1.2638	1.1045	1.3248
	(1.63e-06)	(2.2362e-06)	(4.4443e-06)	(2.7727e-06)	(2.6828e-06)	(0.14286)
1.0 3.0%	3.4906	3.4720	3.4348	1.5319	1.8042	0.90642
	(8.0038e-09)	(1.2728e-08)	(3.311e-08)	(8.7395e-08)	(3.7908e-08)	(0.15809)
1.5 3.0%	5.4727	5.4446	5.3885	1.5938	3.0706	0.56826
	(2.6493e-11)	(1.101e-10)	(2.1414e-09)	(8.1131e-10)	(5.6956e-10)	(0.10693)

Table 4: Huisman-type estimates of the tail-index, for different rates of risk aversion and pure time preference (rows), and for different estimators of the tail-index (columns). No results use equity-weights. All results use FUND.

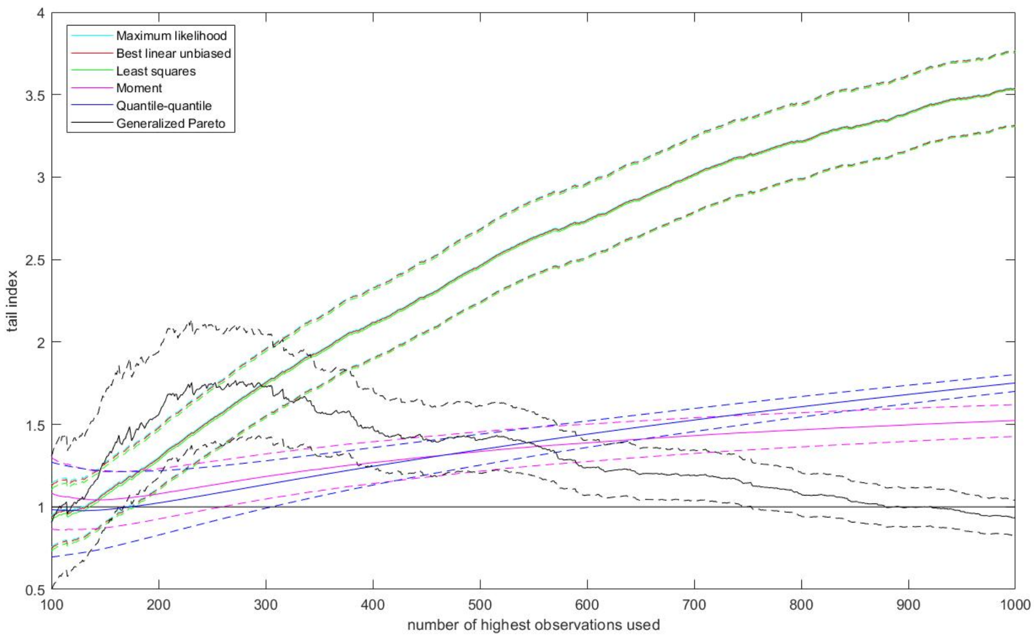

Figure 4: Estimates of the tail-index for FUND, no equity weights.

	ML	BLUE	LS	MM	QQ	GP
DICE 2.5%	16.0946	16.0015	15.8153	-39.2681	16.9784	-0.04886
	(1.2382e-38)	(1.3047e-38)	(1.466e-38)	(301.1745)	(3.7014e-45)	(0.029198)
FUND 2.5%	1.0312	1.0234	1.0079	1.0516	0.88299	1.7445
	(9.8428e-06)	(1.08e-05)	(1.3068e-05)	(1.8199e-05)	(1.3219e-05)	(0.0045129)
PAGE 2.5%	3.4852	3.4653	3.4257	6.3665	3.7482	0.13175
	(6.615e-06)	(6.2179e-06)	(5.4657e-06)	(3.1254e-19)	(3.4169e-08)	(14.3092)
DICE 3.0%	17.2696	17.1696	16.9696	1.1723	18.2353	-0.060949
	(1.6718e-40)	(1.7747e-40)	(2.0178e-40)	(166.6634)	(5.9053e-49)	(0.041988)
FUND 3.0%	1.2723	1.2637	1.2464	1.0485	0.84893	1.6813
	(8.0022e-06)	(9.2109e-06)	(1.241e-05)	(1.221e-05)	(9.7887e-06)	(0.0061745)
PAGE 3.0%	3.573	3.5527	3.5121	14.3872	3.9304	0.10347
	(9.1639e-06)	(8.8766e-06)	(8.364e-06)	(44.4418)	(3.4153e-08)	(0.03559)
DICE 5.0%	23.7391	23.6007	23.324	-4.2772	25.2199	-0.095908
	(2.0356e-58)	(2.1832e-58)	(2.5127e-58)	(228.9773)	(2.1256e-69)	(0.028709)
FUND 5.0%	3.1996	3.1826	3.1488	1.4982	1.6899	0.90946
	(1.7922e-08)	(3.0252e-08)	(1.0431e-07)	(2.4198e-06)	(9.0771e-08)	(0.14391)
PAGE 5.0%	3.657	3.6359	3.5938	26.4057	4.0601	0.078293
	(5.2566e-06)	(5.1115e-06)	(4.8477e-06)	(188.2211)	(1.8641e-08)	(0.035743)

Table 5: Huisman-type estimates of the tail-index, for different models and different constant consumption discount rates (rows), and for different estimators of the tail-index columns). All results use equity-weights.

3.1.5. Consumption discount rates

Table 5 shows the estimated tail-indices for a constant consumption discount rates. Although theoretically inferior to the Ramsey discount rate (Arrow et al., 2013, 2014), constant consumption discount rates are still in use for policy advice by government agencies, particularly those of the USA (on the Social Cost of Carbon, 2010, 2013).

Table 5 confirms the results above. Higher discount rates imply thinner tails, and DICE is more optimistic about the risks of climate change than PAGE, which in turn is more optimistic than FUND.

3.2. Meta-analysis

Table C.7 shows the mean and standard deviation of all estimates in the meta-analysis that use a pure rate of time preference of 1. Taking all estimates at face value leads to a rather high average. These high estimates are discounted by the authors of the studies: The author-weighted average is a factor 25 lower. Adding quality-weights for the studies doubles the social cost of carbon again. The standard deviations show that the range of estimates is rather large.

Figure C.9 shows estimates of the tail-index of the probability density function of the published estimates of the social cost of carbon, taking all estimates at face value. All estimators[5]

[5]Recall that the Generalized Pareto estimator cannot be used on weighted data, and is therefore omitted.

	ML	BLUE	LS	Moment	QQ
no weights	1.2260	1.2136	1.1887	1.3180	1.1850
	(0.3381)	(0.3310)	(0.3214)	(0.0947)	(0.0870)
author weights	1.0766	1.0664	1.0460	1.0017	1.0314
	(0.4572)	(0.4601)	(0.4672)	(0.1037)	(0.0916)
quality weights	1.2241	1.2118	1.1871	1.3047	1.1857
	(0.3499)	(0.3437)	(0.3359)	(0.0933)	(0.0867)

Table 6: Huisman-type estimates of the tail-index, for different weightings of the data, and for different estimators of the tail-index.

for all samples agree that the tail-index is significantly smaller than 2. That is, the tail is thick—the variance does not exist. However, the central estimate of the tail-index is typically larger than 1, although never significantly so. We cannot reject the hypothesis that the tail is fat—we cannot be sure that the mean exists.

Figure C.10 repeats the analysis, now placing lower weight on results that the authors of the estimates themselves de-emphasized. Compared to Figure C.9, we see a downward shift in the estimates of the tail-index. For the moment and quantile-quantile estimators, it is a toss-up whether the tail is fat. For the maximum likelihood, least squares and best linear unbiased estimators, for the 350-500 largest observations, the null hypothesis that the tail-index is 1 or larger is rejected: The tail is fat.

However, as shown in Figure C.10, if we also weigh the estimates by the quality of the study, the tail-index shifts upwards again. The more extreme estimates of the social cost of carbon appear in lower quality papers. With these weights, the tail is thick, but we can neither reject the hypothesis that it is fat nor that it is not fat.

Table 6 provides some clarity. It shows the Huisman et al. (2001) estimators that are independent of tail sample size. The tail-index is larger than one, but not significantly so. Published estimates of the social cost of carbon cannot reject Weitzman's Dismal Theorem.

4. Discussion and Conclusion

Weitzman's Dismal Theorem states that the expectation of the social cost of carbon does not exist. Expected utility maximization can therefore not be used to inform greenhouse gas emission reduction policy. We do not dispute the Dismal Theorem, but it is based on a set of stringent assumptions. Instead, we test whether, if those assumptions are relaxed, the Dismal Theorem still holds. According to the Generalized Central Limit Theorem, the tail of a distribution converges to that of a Pareto distribution, characterised by its tail-index. If the tail-index is smaller than n, the nth moment does not exist. We estimate the tail-index, using a variety of estimators, for Monte Carlo results from three integrated assessment models and for previously published estimates of the social cost of carbon. For the meta-analysis, we cannot reject the hypothesis that the tail-index is smaller than one. For the PAGE model, the tail-index is greater than two, for the DICE much greater than

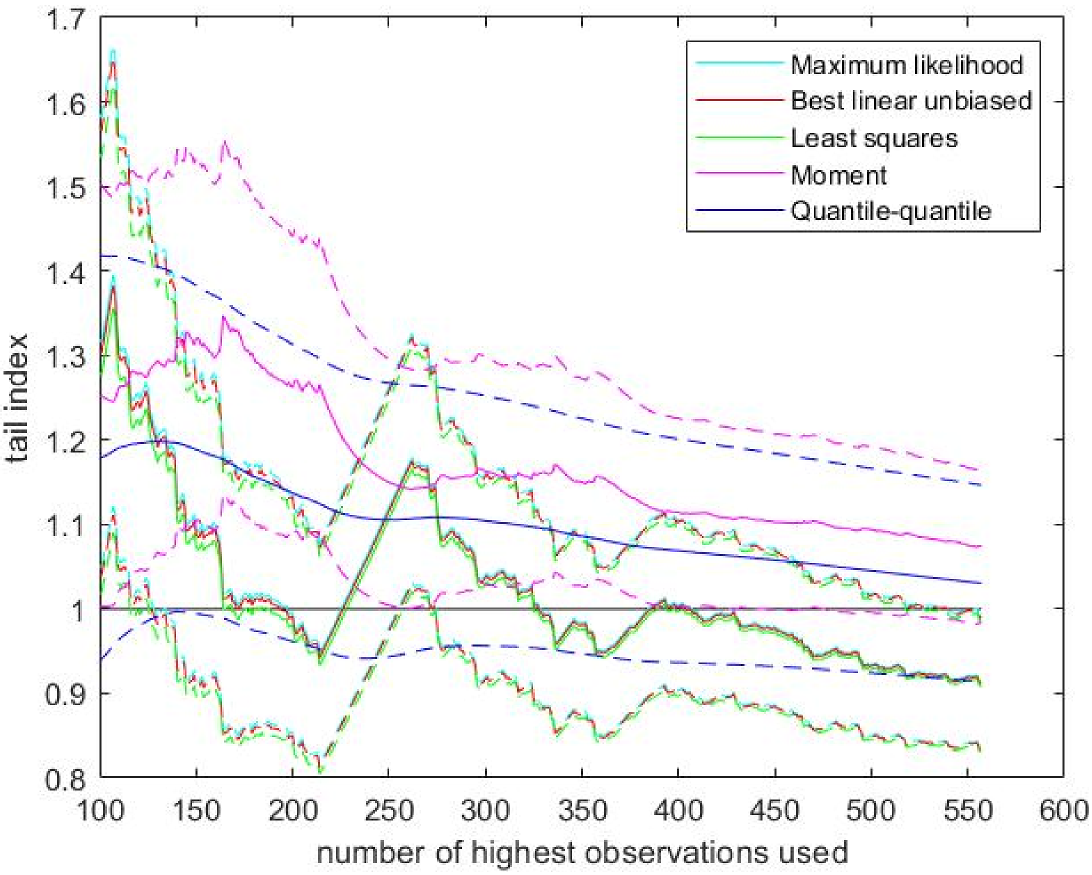

Figure 5: Estimates of the tail-index of the published estimates of the social cost of carbon, using author and quality weights.

two. For FUND, however, the tail-index is smaller than one for a low discount rate, and greater than two for a high discount rate. The difference between DICE and FUND is that the former (latter) has multiplicative (additive) damages, so that utility approaches minus infinity when the damages approach infinity (income) (see Weitzman, 2009b, 2010, for further elaboration).[6] If there is a chance that FUND correctly represents the risks of climate change, we cannot reject the Dismal Theorem. A meta-analysis of published estimates of the social cost of carbon confirms that we cannot exclude a fat tail.

We conclude that the evidence for Weitzman-like dismal theorem results in integrated assessment models is mixed. Some models do not show any signs of fat tails in social cost of carbon estimates, while others do produce fat tailed estimates for some discounting and equity weighting schemes, but not for others. Similarly, we find some evidence for fat tails in a meta-analysis of published social cost of carbon estimates, but those findings are not robust to different quality criteria for our meta analysis. It is clear from the two lines of evidence we investigated that Weitzman's dismal theorem is not just a theoretical curiosity, but might well be found in less stylized and more complex modelling exercises as well.

Our results indicate strongly that more work needs to be done to fully understand when, how and why fat tails emerge in integrated assessment models. Having established a test for dismality, a systematic sensitivity analysis should reveal what exact assumptions drive the tail-index below one. We here find that the rates of time preference, risk aversion and inequity aversion matter, but also that some models do have a fat tail and others do not. The analytical literature shows that the Dismal Theorem is particular to CRRA utility Millner (2013), to no climate policy (Millner, 2013), and to partial equilibrium (Horowitz and Lange, 2014). That is the case for Weitzman's assumptions. Is it true, too, for integrated assessment models? Existing literature gives conflicting answers for a global social planner (Botzen and van den Bergh, 2012, Hwang et al., 2013, Bistline, 2015, Kelly and Tan, 2015, Hwang et al., 2016, Berger et al., 2017, Hwang et al., 2017, 2019, Ikefuji et al., 2020), but there is agreement that catastrophic risk increases cooperation between sovereign nations (Barrett and Dannenberg, 2012, Dellink et al., 2013). Perhaps most importantly, the Dismal Theorem says we cannot use expected cost-benefit analysis to inform climate policy. Unable to reject the Dismal Theorem, alternative criteria to set the desirable intensity of greenhouse gas emission reduction need to be found. This discussion has started (Anthoff and Tol, 2014, Grechuk and Zabarankin, 2014, Aurland-Bredesen, 2020), but not concluded. The same is true for the wider policy implications of the Dismal Theorem (Nordhaus, 2011, Pindyck, 2011, Weitzman, 2011, Nordhaus, 2012, Tol, 2012, Convery and Wagner, 2015, Martin and Pindyck, 2015).

For now, great care must be taken when results from integrated assessment models are used in the policy space. Minimally tests like the ones we present in this paper should routinely be run on model results before they are used to guide policy. If these tests indicate that a given moment of a distribution does not exist, policy makers should refrain from using model

[6] PAGE has additive damages, but triangular distributions for its parameters.

outputs that reflect such a moment. The results in this paper suggest that this problem is more prevalent for low discounting rates, and that for example the discount rates chosen for recent official US government social cost of carbon estimates do not lead to the kind of fat tail that we identified as problematic in this paper.

References

I. B. Aban and M. M. Meerschaert. Generalized least-squares estimators for the thickness of heavy tails. *Journal of Statistical Planning and Inference*, 119(2):341 – 352, 2004. ISSN 0378-3758. URL http://www.sciencedirect.com/science/article/pii/S0378375802004196.

D. Anthoff and J. Emmerling. Inequality and the social cost of carbon. *Journal of the Association of Environmental and Resource Economists*, 6(2):243–273, 2019.

D. Anthoff and R. Tol. On international equity weights and national decision making on climate change. *Journal of Environmental Economics and Management*, 60(1):14–20, 2010. doi: 10.1016/j.jeem.2010.04.002.

D. Anthoff and R. Tol. Climate policy under fat-tailed risk: an application of FUND. *Annals of Operations Research*, 220(1):223–237, September 2014. URL https://ideas.repec.org/a/spr/annopr/v220y2014i1p223-23710.1007-s10479-013-1343-2.html.

D. Anthoff, C. Hepburn, and R. Tol. Equity weighting and the marginal damage costs of climate change. *Ecological Economics*, 68(3):836–849, 2009a. doi: 10.1016/j.ecolecon.2008.06.017.

D. Anthoff, R. Tol, and G. Yohe. Risk aversion, time preference, and the social cost of carbon. *Environmental Research Letters*, 4(2), 2009b. doi: 10.1088/1748-9326/4/2/024002.

D. Anthoff, F. Estrada, and R. Tol. Shutting down the thermohaline circulation. *American Economic Review*, 106(5):602–606, 2016. doi: 10.1257/aer.p20161102.

K. Arrow, M. Cropper, C. Gollier, B. Groom, G. Heal, R. Newell, W. Nordhaus, R. Pindyck, W. Pizer, P. Portney, T. Sterner, R. Tol, and M. Weitzman. Determining benefits and costs for future generations. *Science*, 341(6144):349–350, 2013. doi: 10.1126/science.1235665.

K. Arrow, M. Cropper, C. Gollier, B. Groom, G. Heal, R. Newell, W. Nordhaus, R. Pindyck, W. Pizer, P. Portney, T. Sterner, R. Tol, and M. Weitzman. Should governments use a declining discount rate in project analysis? *Review of Environmental Economics and Policy*, 8(2):145–163, 2014. doi: 10.1093/reep/reu008.

K. Aurland-Bredesen. The benefit-cost ratio as a decision criteria when managing catastrophes. *Environmental and Resource Economics*, 2020. doi: 10.1007/s10640-020-00498-x.

S. Barrett and A. Dannenberg. Climate negotiations under scientific uncertainty. *Proceedings of the National Academy of Sciences of the United States of America*, 109(43):17372–17376, 2012. doi: 10.1073/pnas.1208417109.

L. Berger, J. Emmerling, and M. Tavoni. Managing catastrophic climate risks under model uncertainty aversion. *Management Science*, pages 749–765, 2017. doi: 10.1287/mnsc.2015.2365.

J. Bistline. Fat-tailed uncertainty, learning, and climate policy. *Climate Change Economics*, 6(2), 2015. doi: 10.1142/S2010007815500098.

W. Botzen and J. van den Bergh. How sensitive is nordhaus to weitzman? climate policy in DICE with an alternative damage function. *Economics Letters*, 117(1):372–374, 2012. doi: 10.1016/j.econlet.2012.05.032.

L. Bretschger and A. Pattakou. As bad as it gets: How climate damage functions affect growth and the social cost of carbon. *Environmental and Resource Economics*, 72(1):5–26, 2019. doi: 10.1007/s10640-018-0219-y.

M. Brito and A. C. M. Freitas. Limiting behaviour of a geometric-type estimator for tail indices. *Insurance: Mathematics and Economics*, 33(2):211 – 226, 2003. ISSN 0167-6687. URL http://www.sciencedirect.com/science/article/pii/S0167668703001355.

M. Budolfson, F. Dennig, M. Fleurbaey, A. Siebert, and R. H. Socolow. The comparative importance for optimal climate policy of discounting, inequalities and catastrophes. *Climatic Change*, 145(3):481–494, 2017. URL https://doi.org/10.1007/s10584-017-2094-x.

S. Cato. From the st. petersburg paradox to the dismal theorem. *Environment and Development Economics*, 2020. doi: 10.1017/S1355770X20000121.

F. Convery and G. Wagner. Managing uncertain climates: Some guidance for policy makers and researchers. *Review of Environmental Economics and Policy*, 9(2):304–320, 2015. doi: 10.1093/reep/rev003.

K. D. Daniel, R. B. Litterman, and G. Wagner. Declining CO_2 price paths. *Proceedings of the National Academy of Sciences*, 116(42):20886–20891, 2019. URL https://www.pnas.org/content/116/42/20886.

K. Dayaratna, R. McKitrick, and P. Michaels. Climate sensitivity, agricultural productivity and the social cost of carbon in FUND. *Environmental Economics and Policy Studies*, 2020. doi: 10.1007/s10018-020-00263-w.

A. L. M. Dekkers, J. H. J. Einmahl, and L. D. Haan. A moment estimator for the index of an extreme-value distribution. *The Annals of Statistics*, 17(4):1833–1855, 12 1989. doi: 10.1214/aos/1176347397. URL https://doi.org/10.1214/aos/1176347397.

R. Dellink, T. Dekker, and J. Ketterer. The fatter the tail, the fatter the climate agreement: Simulating the influence of fat tails in climate change damages on the success of international climate negotiations. *Environmental and Resource Economics*, 56(2):277–305, 2013. doi: 10.1007/s10640-013-9642-2.

T. Ekholm. Climatic cost-benefit analysis under uncertainty and learning on climate sensitivity and damages. *Ecological Economics*, 154:99–106, 2018. doi: 10.1016/j.ecolecon.2018.07.024.

S. Fankhauser, R. Tol, and D. Pearce. The aggregation of climate change damages: A welfare theoretic approach. *Environmental and Resource Economics*, 10(3):249–266, 1997. doi: 10.1023/A:1026420425961.

T. Faulwasser, R. Nydestedt, C. Kellett, and S. Weller. Towards a FAIR-DICE IAM: Combining DICE and FAIR models. *IFAC-PapersOnLine*, 51(5):126–131, 2018. doi: 10.1016/j.ifacol.2018.06.222.

I. Fedotenkov. A review of more than one hundred pareto-tail index estimators. Technical Report 90072, Munich Personal RePEc Archive, 2018. URL https://mpra.ub.uni-muenchen.de/90072/.

J. Geweke. A note on some limitations of crra utility. *Economics Letters*, 71(3):341 – 345, 2001. ISSN 0165-1765. doi: https://doi.org/10.1016/S0165-1765(01)00391-3. URL http://www.sciencedirect.com/science/article/pii/S0165176501003913.

N. Glanemann, S. N. Willner, and A. Levermann. Paris climate agreement passes the cost-benefit test. *Nature Communications*, 11(1):110, 2020. URL https://doi.org/10.1038/s41467-019-13961-1.

A. Golub and M. Brody. Uncertainty, climate change, and irreversible environmental effects: application of real options to environmental benefit-cost analysis. *Journal of Environmental Studies and Sciences*, 7(4):519–526, 2017. doi: 10.1007/s13412-017-0436-7.

B. Grechuk and M. Zabarankin. Risk averse decision making under catastrophic risk. *European Journal of Operational Research*, 239(1):166–176, 2014. doi: 10.1016/j.ejor.2014.04.042.

C. Guivarch and A. Pottier. Climate damage on production or on growth: What impact on the social cost of carbon? *Environmental Modeling and Assessment*, 23(2):117–130, 2018. doi: 10.1007/s10666-017-9572-4.

S. Hafeez, S. Weller, and C. M. Kellett. Impact of climate model parametric uncertainty in an MPC implementation of the DICE integrated assessment model. *IFAC-PapersOnLine*, 50(1):959–965, 2017. doi: 10.1016/j.ifacol.2017.08.169.

M. Hänsel and M. Quaas. Intertemporal distribution, sufficiency, and the social cost of carbon. *Ecological Economics*, 146:520–535, 2018. doi: 10.1016/j.ecolecon.2017.11.024.

B. M. Hill. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics*, 3(5):1163–1174, 1975. ISSN 00905364. URL http://www.jstor.org/stable/2958370.

C. Hope and K. Schaefer. Economic impacts of carbon dioxide and methane released from thawing permafrost. *Nature Climate Change*, 6(1):56–59, 2016. doi: 10.1038/nclimate2807.

J. Horowitz and A. Lange. Cost–benefit analysis under uncertainty — a note on weitzman's dismal theorem. *Energy Economics*, 42:201 – 203, 2014. URL http://www.sciencedirect.com/science/article/pii/S014098831300296X.

R. Huisman, K. G. Koedijk, C. J. M. Kool, and F. Palm. Tail-index estimates in small samples. *Journal of Business & Economic Statistics*, 19(2):208–216, 2001. URL http://www.jstor.org/stable/1392164.

I. Hwang, F. Reynès, and R. Tol. Climate policy under fat-tailed risk: An application of DICE. *Environmental and Resource Economics*, 56(3):415–436, 2013. doi: 10.1007/s10640-013-9654-y.

I. Hwang, R. Tol, and M. Hofkes. Fat-tailed risk about climate change and climate policy. *Energy Policy*, 89:25–35, 2016. doi: 10.1016/j.enpol.2015.11.012.

I. Hwang, F. Reynès, and R. Tol. The effect of learning on climate policy under fat-tailed risk. *Resource and Energy Economics*, 48:1–18, 2017. doi: 10.1016/j.reseneeco.2017.01.001.

I. Hwang, R. Tol, and M. Hofkes. Active learning and optimal climate policy. *Environmental and Resource Economics*, 73(4):1237–1264, 2019. doi: 10.1007/s10640-018-0297-x.

M. Ikefuji, R. Laeven, J. Magnus, and C. Muris. Expected utility and catastrophic risk in a stochastic economy–climate model. *Journal of Econometrics*, 214(1):110–129, 2020. doi: 10.1016/j.jeconom.2019.05.007.

D. Kelly and Z. Tan. Learning and climate feedbacks: Optimal climate insurance and fat tails. *Journal of Environmental Economics and Management*, 72:98–122, 2015. doi: 10.1016/j.jeem.2015.05.001.

M. Kotchen. Which social cost of carbon? a theoretical perspective. *Journal of the Association of Environmental and Resource Economists*, 5(3):673–694, 2018. doi: 10.1086/697241.

M. Kratz and S. I. Resnick. The qq-estimator and heavy tails. *Communications in Statistics. Stochastic Models*, 12(4):699–724, 1996. URL https://doi.org/10.1080/15326349608807407.

I. Martin and R. Pindyck. Averting catastrophes: The strange economics of Scylla and Charybdis. *American Economic Review*, 105(10):2947–2985, 2015. doi: 10.1257/aer.20140806.

A. Millner. On welfare frameworks and catastrophic climate risks. *Journal of Environmental Economics and Management*, 65(2):310 – 325, 2013. URL http://www.sciencedirect.com/science/article/pii/S0095069612001209.

F. Moore, U. Baldos, T. Hertel, and D. Diaz. New science of climate change impacts on agriculture implies higher social cost of carbon. *Nature Communications*, 8(1), 2017. doi: 10.1038/s41467-017-01792-x.

F. Moore, J. Rising, N. Lollo, C. Springer, V. Vasquez, A. Dolginow, C. Hope, and D. Anthoff. Mimi-page, an open-source implementation of the page09 integrated assessment model. *Scientific Data*, 5, 2018. doi: 10.1038/sdata.2018.187.

E. National Academies of Sciences and Medicine. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide*. The National Academies Press, Washington, DC, 2017. ISBN 978-0-309-45420-9. URL https://www.nap.edu/catalog/24651/valuing-climate-damages-updating-estimation-of-the-social-cost-of.

W. Nordhaus. The economics of tail events with an application to climate change. *Review of Environmental Economics and Policy*, 5(2):240–257, 2011. doi: 10.1093/reep/rer004.

W. Nordhaus. Economic policy in the face of severe tail events. *Journal of Public Economic Theory*, 14(2):197–219, 2012. doi: 10.1111/j.1467-9779.2011.01544.x.

W. Nordhaus. Evolution of modeling of the economics of global warming: changes in the dice model, 1992–2017. *Climatic Change*, 148(4):623–640, 2018. doi: 10.1007/s10584-018-2218-y.

W. D. Nordhaus. How fast should we graze the global commons? *American Economic Review*, 72(2):242–246, 1982.

W. D. Nordhaus. An optimal transition path for controlling greenhouse gases. *Science*, 258:1315–1319, 1992.

W. D. Nordhaus. Rolling the 'DICE': An optimal transition path for controlling greenhouse gases. *Resource and Energy Economics*, 15(1):27–50, 1993.

W. D. Nordhaus. Climate clubs: Overcoming free-riding in international climate policy. *American Economic Review*, 105(4):1339–70, April 2015. URL http://www.aeaweb.org/articles?id=10.1257/aer.15000001.

S. Okullo. Determining the social cost of carbon: Under damage and climate sensitivity uncertainty. *Environmental and Resource Economics*, 75(1):79–103, 2020. doi: 10.1007/s10640-019-00389-w.

I. W. G. on the Social Cost of Carbon. Technical support document: Social cost of carbon for regulatory impact analysis under executive order 12866. Report, United States Government, 2010.

I. W. G. on the Social Cost of Carbon. Technical support document: Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866. Report, United States Government, 2013.

J. Paul Kelleher and G. Wagner. Ramsey discounting calls for subtracting climate damages from economic growth rates. *Applied Economics Letters*, 26(1):79–82, 2019. doi: 10.1080/13504851.2018.1438581.

A. Pigou. *The Economics of Welfare*. Macmillan, 1920.

R. Pindyck. Fat tails, thin tails, and climate change policy. *Review of Environmental Economics and Policy*, 5(2):258–274, 2011. doi: 10.1093/reep/rer005.

R. Pindyck. Coase lecture—taxes, targets and the social cost of carbon. *Economica*, 84(335):345–364, 2017. doi: 10.1111/ecca.12243.

R. Pindyck. The social cost of carbon revisited. *Journal of Environmental Economics and Management*, 94: 140–160, 2019. doi: 10.1016/j.jeem.2019.02.003.

E. Plambeck, C. Hope, and J. Anderson. The page95 model: Integrating the science and economics of global warming. *Energy Economics*, 19(1):77–101, 1997. doi: 10.1016/s0140-9883(96)01008-0.

K. Ricke, L. Drouet, K. Caldeira, and M. Tavoni. Country-level social cost of carbon. *Nature Climate Change*, 8(10):895–900, 2018. doi: 10.1038/s41558-018-0282-y.

J. Schultze and J. Steinebach. On least squares estimates of an exponential tail coefficient. *Statistics Risk Modeling*, 14(4):353–372, 1996.

N. Scovronick, M. Budolfson, F. Dennig, M. Fleurbaey, A. Siebert, R. Socolow, D. Spears, and F. Wagner. Impact of population growth and population ethics on climate change mitigation policy. *Proceedings of the National Academy of Sciences of the United States of America*, 114(46):12338–12343, 2017. doi: 10.1073/pnas.1618308114.

R. Tol. Spatial and temporal efficiency in climate policy: Applications of fund. *Environmental and Resource Economics*, 14(1):33–49, 1999. doi: 10.1023/A:1008314205375.

R. Tol. Estimates of the damage costs of climate change: Part 1: Benchmark estimates. *Environmental and Resource Economics*, 21(1):47–73, 2002a. doi: 10.1023/A:1014500930521.

R. Tol. Estimates of the damage costs of climate change: Part ii. dynamic estimates. *Environmental and Resource Economics*, 21(2):135–160, 2002b. doi: 10.1023/A:1014539414591.

R. Tol. Leviathan carbon taxes in the short run: A letter. *Climatic Change*, 114(2):409–415, 2012. doi: 10.1007/s10584-012-0544-z.

R. S. Tol. Kernel density decomposition with an application to the social cost of carbon. Working Paper Series 0720, Department of Economics, University of Sussex Business School, 2020. URL https://ideas.repec.org/p/sus/susewp/0720.html.

R. S. J. Tol. The Economic Impacts of Climate Change. *Review of Environmental Economics and Policy*, 12(1):4–25, 01 2018. URL https://doi.org/10.1093/reep/rex027.

R. S. J. Tol. A social cost of carbon for (almost) every country. *Energy Economics*, 83:555–566, 2019. doi: 10.1016/j.eneco.2019.07.006.

Y. M. Tripathi, S. Kumar, and C. Petropoulos. Improved estimators for parameters of a pareto distribution with a restricted scale. *Statistical Methodology*, 18:1 – 13, 2014. ISSN 1572-3127. URL http://www.sciencedirect.com/science/article/pii/S1572312713000749.

F. van der Ploeg and A. de Zeeuw. Pricing carbon and adjusting capital to fend off climate catastrophes. *Environmental and Resource Economics*, 72(1):29–50, 2019. doi: 10.1007/s10640-018-0231-2.

M. Weitzman. What is the "damages function" for global warming-and what difference might it make? *Climate Change Economics*, 1(1):57–69, 2010. doi: 10.1142/S2010007810000042.

M. Weitzman. Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*, 5(2):275–292, 2011. doi: 10.1093/reep/rer006.

M. L. Weitzman. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, 91(1):1–19, 2009a. doi: 10.1162/rest.91.1.1.

M. L. Weitzman. Additive damages, fat-tailed climate dynamics, and uncertain discounting. *Economics - The Open-Access, Open-Assessment E-Journal*, 3:1–29, 2009b. doi: 10.5018/economics-ejourna.

P. Yang, Y.-F. Yao, Z. Mi, Y.-F. Cao, H. Liao, B.-Y. Yu, Q.-M. Liang, D. Coffman, and Y.-M. Wei. Social cost of carbon under shared socioeconomic pathways. *Global Environmental Change*, 53:225–232, 2018. doi: 10.1016/j.gloenvcha.2018.10.001.

Z. Zhen and L. Tian. The impact of climate damage function on the social cost of carbon and economic growth rate. *Mitigation and Adaptation Strategies for Global Change*, 2019. doi: 10.1007/s11027-019-09908-y.

Z. Zhen, L. Tian, and Q. Ye. A simple estimate for the social cost of carbon. volume 152, pages 768–773, 2018. doi: 10.1016/j.egypro.2018.09.243.

Appendix A. Tail-index estimators

Let $X_1, X_2, ..., X_n$ denote a series of observations and $X_{(1)}, X_{(2)}, X_{(n)}$ its order statistics. The Pareto distribution is defined as

$$F(X) = 1 - \left(\frac{\beta}{X}\right)^{\alpha} \tag{A.1}$$

for $X \geq \beta > 0$, with density

$$f(X) = \frac{\alpha\beta^{\alpha}}{X^{\alpha+1}} \tag{A.2}$$

where α is the tail-index and $\gamma = \alpha^{-1}$ its inverse.

The Generalized Central Limit Theorem has that the sum of independent, identically and symmetrically distributed random variables converges to a distribution whose tail is Pareto, with $0 < \alpha \leq 2$, with $\alpha = 2$ if the variance is finite.

A useful statistic is

$$M_l(k) = \frac{1}{\sum_{i=0}^{k-1} w_i} \sum_{i=0}^{k-1} w_i \left(\ln X_{(n-i)} - \ln X_{(n-k)}\right)^l \tag{A.3}$$

where $w_i = 1$ (for now). M_l is the l^{th} non-central moment of the slope between the k^{th}-largest observation and larger ones.

The maximum likelihood estimator of the tail-index (Hill, 1975) is

$$\hat{\alpha}^H(k) = M_1(k)^{-1} \tag{A.4}$$

Its asymptotic distribution is Normal with standard error

$$\hat{\sigma}_{\alpha}^H(k) = \frac{k}{(k-1)\sqrt{k-2}}\hat{\alpha}^H(k) \tag{A.5}$$

Aban and Meerschaert (2004) show that

$$\hat{\alpha}^{AM}(k) = \frac{k-1}{k}\hat{\alpha}^H(k) \tag{A.6}$$

is the best linear unbiased estimator of the inverse tail-index,[7] and its uniform minimum variance unbiased estimator. Its asymptotic standard error is

$$\hat{\sigma}_{\alpha}^{AM}(k) = \frac{1}{\sqrt{k-2}} \hat{\alpha}^{H}(k) \tag{A.7}$$

Tripathi et al. (2014) show that

$$\hat{\alpha}^{T}(k) = \frac{k-3}{k} \hat{\alpha}^{H}(k) \tag{A.8}$$

is the least squares estimator of the tail-index. Its asymptotic standard error follows.

Dekkers et al. (1989) propose a moment estimator

$$\left(\hat{\alpha}^{D}(k)\right)^{-1} = M_1(k) + 1 - \frac{1}{2}\left(1 - \frac{M_1(k)^2}{M_2(k)}\right)^{-1} \tag{A.9}$$

Its asymptotic standard error is $\hat{\alpha}^{D}(k)/\sqrt{k}$.

This estimator is unstable if $M_1(k)^2 \approx M_2(k)$ or $\mathbb{V}\text{ar}\left(\ln X_{(n-i)} - \ln X_{(n-k)}\right) \approx 0$—that is, if the order statistics decline exponentially, as in the Normal distribution so that $\alpha \to \infty$. Taking the same ratio for the theoretical moments of the Pareto distribution $\frac{\mu_1^2}{\mu_2} = \frac{\alpha(\alpha-2)}{1+\alpha(\alpha-2)}$ reveals that the Dekkers estimators becomes unstable for thinly tailed distributions.

The above estimators all assume that estimates are of equal quality. In fact, in the meta-analysis, papers present some estimates as central and other estimates as sensitivity analyses or replications of previous studies. Some papers are better than others. The estimators shown above all depend on the M_l statistics, which are readily generalized to weighted data—and indeed already are in Equation (A.3) for $w_i \neq 1$ and $\sum_i w_i = k$.

One problem with the above estimators is that they work well if the right tail of the distribution is exactly Pareto, but not so well if the tail is approximately so. Fedotenkov (2018) reviews many of the suggested solutions. Kratz and Resnick (1996) suggest to run a regression of the natural logarithm of the k largest observations on the natural logarithm of their order, for the Pareto distribution describes Zipf's Law. Schultze and Steinebach (1996) argued that dependent and independent variables should be switched. Brito and Freitas (2003) show that the geometric mean of these two estimators performs better. The asymptotic standard error of these estimators is the estimate times the square root of two divided by the square root of k. This is readily generalized to weighted least squares.

Besides generic deviations from the Pareto distribution, we also consider a particular deviation: The Generalized Pareto Distribution. Its parameters are estimated by *numerical* maximum likelihood for the k largest observations, using the GPFIT function from Matlab's

[7] Hill had already shown that it is unbiased

Statistics toolbox. This is not readily generalized to weighted observations, so we do not apply this to the meta-analysis. Note that, like the moments estimator, GPFIT becomes unstable as the tail-index grows large since the shape parameter is its inverse.

The above estimators all depend on k. Although there is some guidance on how to select k (see Fedotenkov, 2018), we instead show results for a range. We further apply the Huisman estimator.

Huisman et al. (2001) argue for a two-stage estimation: First estimate the tail-index for a range of k, and then regress the estimated tail-indices on k. The intercept from that regression is the Huisman estimator for the tail-index. The regression on k removes its influence, at least to a first-order approximation.

While Huisman *et al.* use the Hill estimator, we here apply the same method to all estimators.

As the variance of the Hill estimator is proportional to k, Huisman et al. (2001) use weighted least squares with weights $\sqrt{k}$. The Hill estimators for different k are correlated because the same data are used.[8] The covariance matrix is therefore

$$Cov(\beta) = (Z'W'WZ)^{-1}Z'W'A\Sigma A'W'WZ(Z'W'WZ)^{-1} \tag{A.10}$$

For $A\Sigma A' = 1$, this is the covariance matrix for a weighted least squares regression, where

$$Z = \begin{pmatrix} 1 & 1 \\ 1 & 2 \\ \dots & \dots \\ 1 & k \\ \dots & \dots \\ 1 & K \end{pmatrix}$$

and

$$W = \begin{pmatrix} 1 & 1 \\ 1 & \sqrt{(2)} \\ \dots & \dots \\ 1 & \sqrt{k} \\ \dots & \dots \\ 1 & \sqrt{K} \end{pmatrix}$$

[8] Anthoff and Tol (2014) did not correct for this.

Models			Meta		
DICE	FUND	PAGE	No	Author	Quality
70	16	100	983	39	82
(30)	(101)	(256)	(3171)	(479)	(765)

Table C.7: Mean and standard deviation of the social cost of carbon (in \$/tC) for the three integrated assessment models, for a pure rate of time preference of 1% per year and a rate of risk aversion of 1, and for the meta-analysis with no weights, author weights and quality weights, for a pure rate of time preference of 1% per year.

The term $A\Sigma A'$ corrects for correlated observations.

$$A = \begin{pmatrix} 0 & ... & 0 & 0 & 0 & -1 & 1 \\ 0 & ... & 0 & 0 & -1 & 1/2 & 1/2 \\ 0 & ... & 0 & -1 & 1/3 & 1/3 & 1/3 \\ . & . & . & . & . & . & . \\ -1 & ... & 1/K & 1/K & 1/K & 1/K & 1/K \end{pmatrix}$$

The elements of Σ are

$$\Sigma(i,j) = \begin{cases} \frac{i/N(1-j/N)}{\hat{\alpha}(i)\hat{\alpha}(j)N} \left(\ln\left(1-\frac{i}{N}\right)^{-1/\hat{\alpha}(i)}\right)^{\hat{\alpha}(i)+1} \left(\ln\left(1-\frac{j}{N}\right)^{-1/\hat{\alpha}(j)}\right)^{\hat{\alpha}(j)+1} & \text{if } i > j \\ 0 & \text{if } i \leq j \end{cases}$$

Note that $\Sigma(i,j)$ is a complex number for $\hat{\alpha}(i) < 0$. In that case, we set $A\Sigma A' = 1$.

The Matlab code is available at GitHub.

Appendix B. New estimates of the social cost of carbon

The previous meta-analysis of the social cost of carbon (Tol, 2018) was extended with estimates reported in Anthoff and Emmerling (2019), Bretschger and Pattakou (2019), Budolfson et al. (2017), Daniel et al. (2019), Dayaratna et al. (2020), Ekholm (2018), Faulwasser et al. (2018), Golub and Brody (2017), Guivarch and Pottier (2018), Hafeez et al. (2017), Hänsel and Quaas (2018), Kotchen (2018), Moore et al. (2017), Nordhaus (2015), Okullo (2020), Ricke et al. (2018), Scovronick et al. (2017), Tol (2019), Yang et al. (2018) and Zhen et al. (2018). The Budolfson and Faulwasser estimates were digitized from graphs.

Glanemann et al. (2020) do not report a carbon tax, Zhen and Tian (2019) report the *relative* social cost of carbon, Paul Kelleher and Wagner (2019) *relative changes* in the social cost of carbon, van der Ploeg and de Zeeuw (2019) the *steady state* social cost of carbon, and Pindyck (2017, 2019) the *average* social cost of carbon.

Appendix C. Additional results

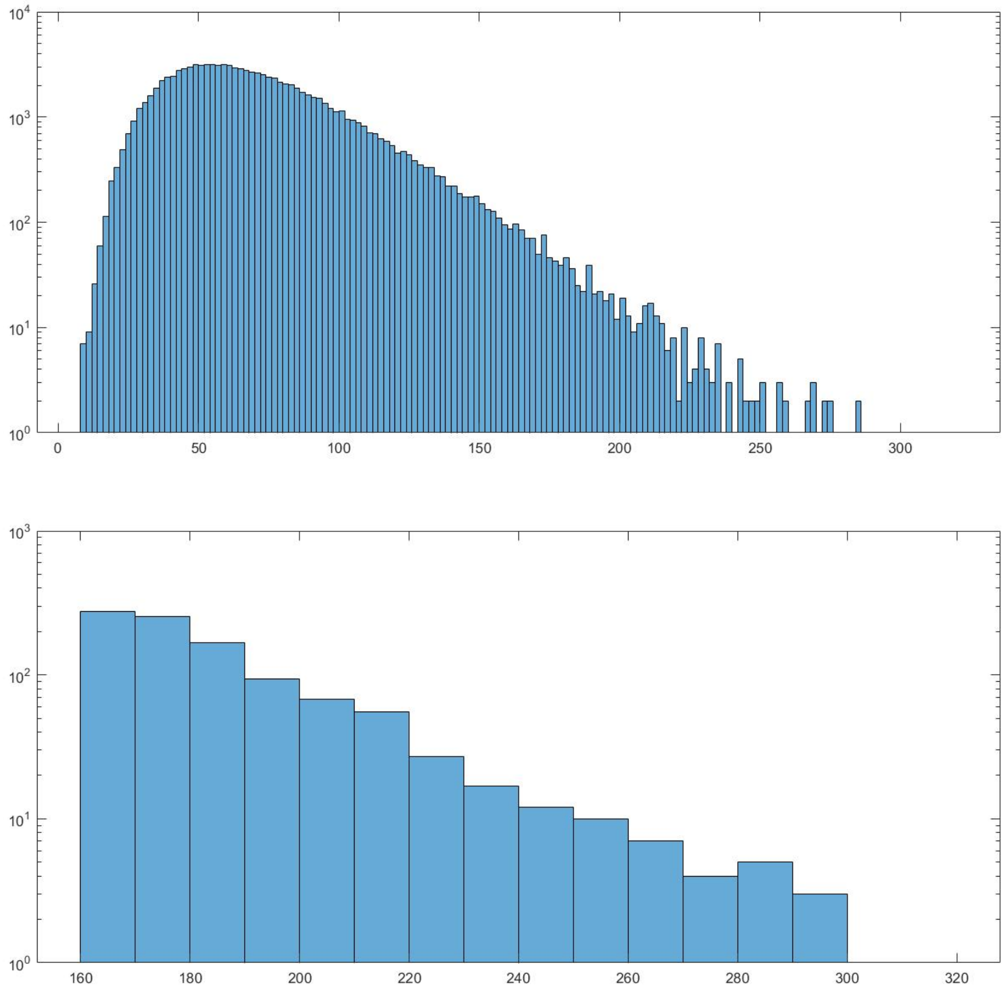

Figure C.6: Histogram of estimates of the social cost of carbon (in \$/tC) by DICE for a rate of risk aversion of 1 and a pure rate of time preference of 1%. The top panel shows all 100,000 draws, the bottom panel the 1000 largest.

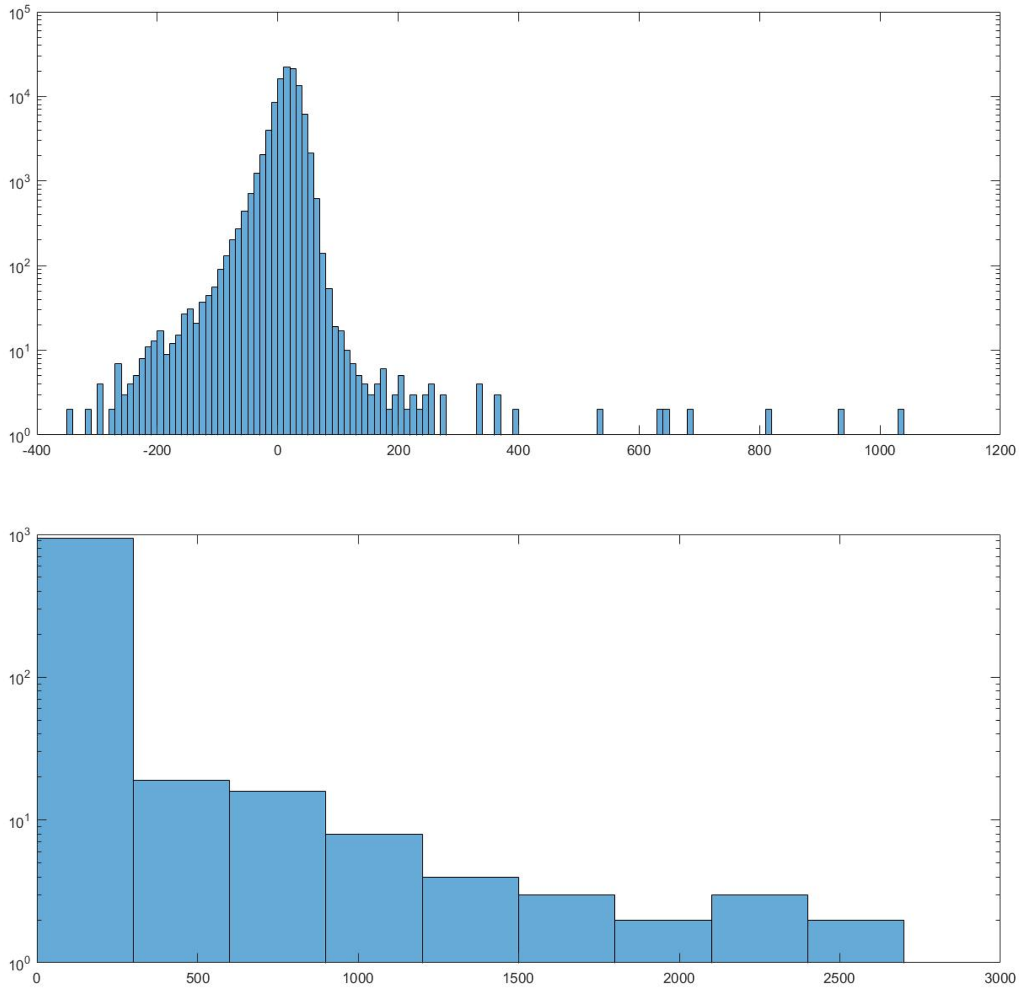

Figure C.7: Histogram of estimates of the social cost of carbon by FUND using equity weights for a rate of risk aversion of 1 and a pure rate of time preference of 1%. The top panel shows all 100,000 draws, the bottom panel the 1000 largest.

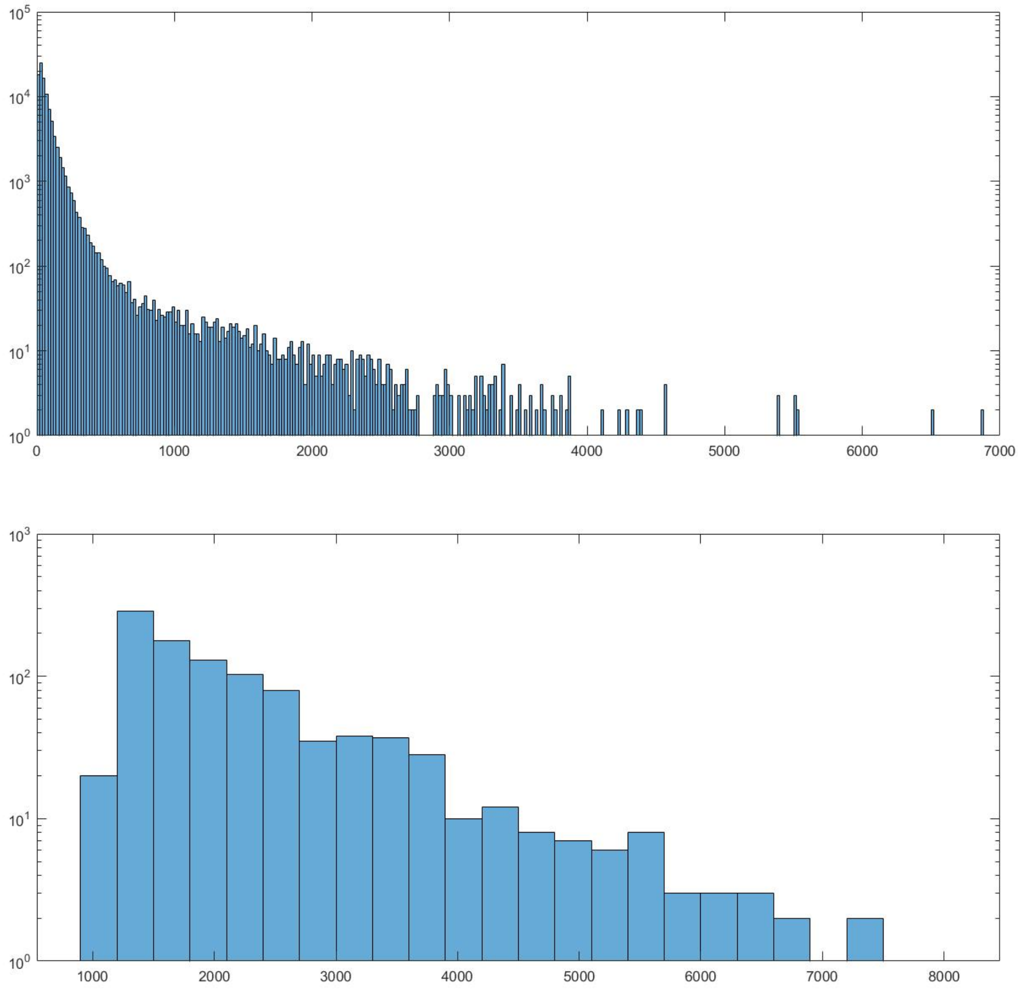

Figure C.8: Histogram of estimates of the social cost of carbon by PAGE using equity weights for a rate of risk aversion of 1 and a pure rate of time preference of 1%. The top panel shows all 100,000 draws, the bottom panel the 1000 largest.

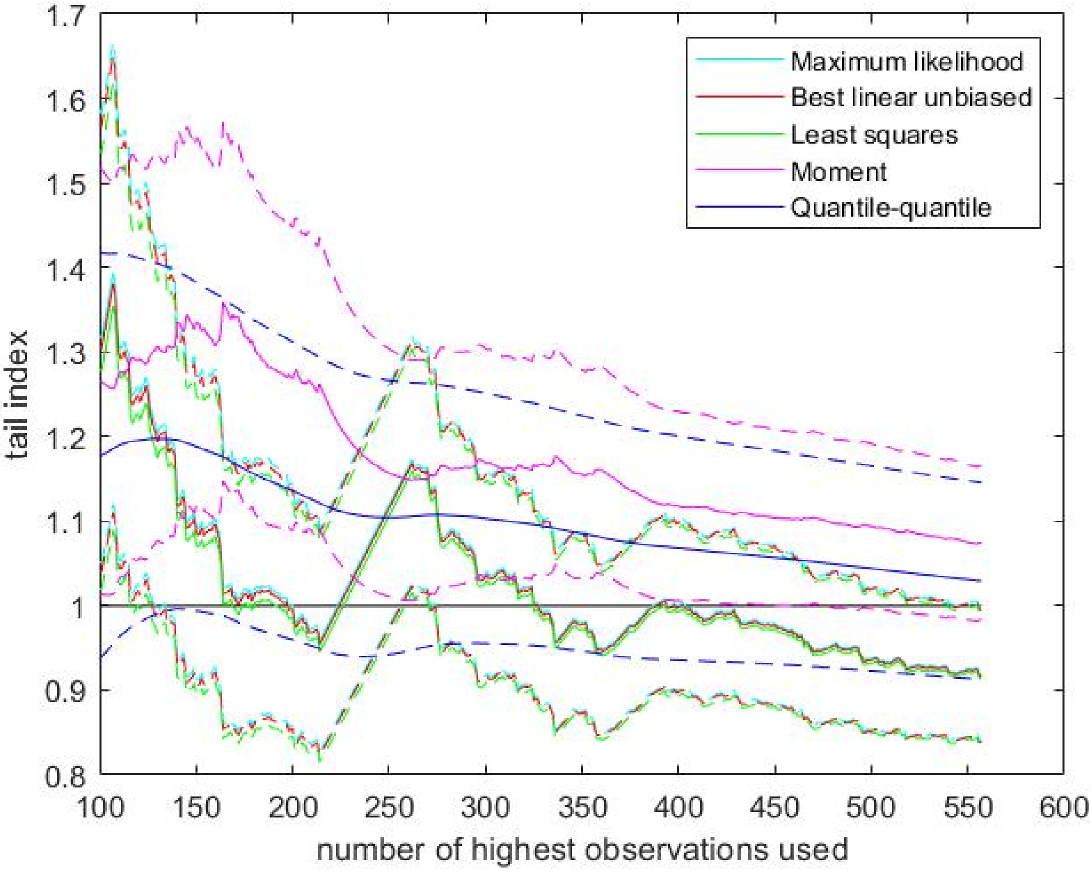

Figure C.9: Estimates of the tail-index of the published estimates of the social cost of carbon.

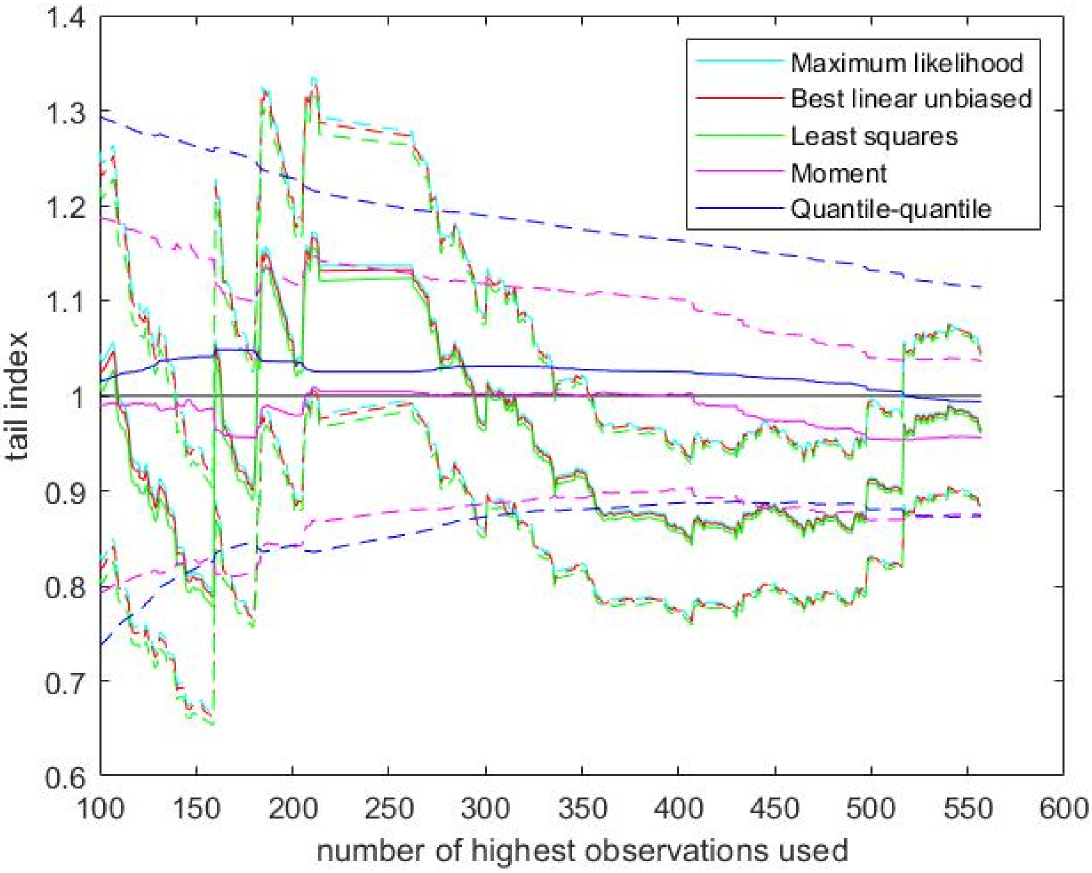

Figure C.10: Estimates of the tail-index of the published estimates of the social cost of carbon, using author weights.