STOCKHOLM
ENVIRONMENT
INSTITUTE

Climate damages in the FUND model: A disaggregated analysis

Frank Ackerman and Charles Munitz

March 2011

Stockholm Environment Institute

Working Paper WP-US-1105

Abstract

The FUND model of climate economics, developed by Richard Tol and David Anthoff, is widely used in research and policy-making. It was one of three models used by the U.S. government's Interagency Working Group on the Social Cost of Carbon in 2009, which estimated the SCC – the cost of incremental damages from greenhouse gas emissions – at $21 per ton of CO_2. This paper presents a disaggregation of the damage estimates in FUND, followed by a more detailed examination of agricultural damages in particular. FUND's own central estimate of the SCC (under the Interagency Working Group assumptions) is $6, the sum of an estimated net benefit in agriculture, a net cost in heating and cooling, and very small net costs in all other areas. We identify flaws in FUND's equations, and find that FUND relies on outdated research that overestimates agricultural benefits from climate change.

For more information about this document,
contact Frank Ackerman at frank.ackerman@sei-us.org.

Introduction

The FUND model of climate economics, developed by Richard Tol and David Anthoff, is widely used, both in research and in the development of policy proposals. It was one of three models used by the U.S. government's Interagency Working Group on the Social Cost of Carbon in 2009 (Interagency Working Group 2010). The Working Group's "central estimate"[1] of the social cost of carbon (SCC), i.e. the monetary value of the incremental damages from greenhouse gas emissions, was $21 per ton of CO2.

FUND differs from the other two models used by the Interagency Working Group, DICE and PAGE, in at least two important respects. First, it produces the lowest central estimate of the SCC, $6, compared with $30 in PAGE and $28 in DICE. (Here and throughout, SCC estimates are in 2007 dollars per ton of CO2.) Second, FUND is far more complex than the other models, with, among other features, 15 major categories and additional subcategories of climate damages, each based on a separate analysis and estimated for each of 16 regions of the world. Many of the constants defining these damages, as well as those used in other aspects of FUND, are modeled as Monte Carlo parameters, often with means and standard deviations specified separately for each region. As a consequence of this level of detail and complexity, it seems likely that many economists and policy analysts who use FUND results are unaware of the contribution of individual features of FUND to the final outcomes.

This paper presents a disaggregation of the damage estimates in FUND, followed by a more detailed examination of agricultural damages in particular. It then raises three issues about the modeling of agricultural damages in FUND, reviews recent literature relevant to agricultural damages, and recommends changes in FUND.

Methodology

The analysis described here begins with the Working Group's modified version of FUND.[2] Software switches were then installed, making it possible to turn off individual damage components while keeping other features of the model unchanged.[3] FUND was then re-run with various categories turned off. Turning off a damage category X produces what might be called the "all-but-X" estimate of the SCC; the impact of X can be defined as the Working Group estimate minus the all-but-X estimate.

This can be done either in the Monte Carlo mode of operation, used in the Working Group analysis and most FUND-based research, or in a best-guess mode, in which each of the Monte Carlo parameters is fixed at its mode or best-guess value. The contrast between Monte Carlo and best-guess results measures the impact of uncertainty as modeled in FUND.

---

[1] "Central estimate," the Working Group's terminology, refers to the estimate of the SCC under assumptions made in the Working Group analysis, including a fixed 3 percent discount rate, other specified inputs, and a set of five scenarios, the results of which are averaged. Results at the 5 percent and 2.5 percent discount rates, also used by the Working Group, are qualitatively similar, and are omitted from this article to simplify the presentation.

[2] Thanks to David Anthoff for providing the FUND files, and for assistance in getting FUND running on our computers. He is, of course, not responsible for any statements about FUND made in this paper.

[3] A technical memo on the software modifications made to FUND for this analysis is available from the authors on request.

Damage calculations play two distinct roles in FUND. First, for market impact categories (i.e., excluding externality valuations), each year's damages are subtracted from the next year's output, reducing the resources available for consumption. Second, for all damage categories, the present value of the future stream of damages is the basis for the calculation of the SCC. In that calculation, the model is run twice with nearly identical patterns of emissions, differing only in an added pulse of emissions in a specific year. The SCC for that year is the present value of the difference between future damages in the two runs, per ton of carbon in the emissions pulse. The Working Group performed this calculation for several years; only the 2010 results are discussed in this paper.

The modifications of FUND described below affect only the second use of the damage estimates (except where noted). That is, in most of the modified FUND runs presented here, there is no change to the damages subtracted from output; the software switches merely remove categories of damages from the calculation of the SCC. This ensures that the trajectories of income, consumption, and investment remain the same as in the Working Group version of FUND, even when portions of the SCC calculation have been turned off.

Results

Comparing FUND and DICE

An initial experiment with FUND demonstrates that the gap between the FUND and DICE "central estimates" of the SCC can be entirely explained by the difference in their treatment of climate damages.[4] In place of FUND's disaggregated analysis, DICE uses a single equation to model damages:

$$(1) \quad output\ net\ of\ damages = \frac{gross\ output}{1+.0028838T^2}$$

Gross output is the output that would have been produced in the absence of climate change, and T is the change in temperature in $^{\circ}$C since 1900 (Nordhaus 2008). When damages are calculated by substituting equation (1) from DICE into the Working Group version of FUND, keeping everything else unchanged, the result is an SCC of \$31 per ton, about 10 percent greater than the DICE value. That is, if the two models agreed on DICE's climate damages, they would roughly agree in their estimates of the SCC.
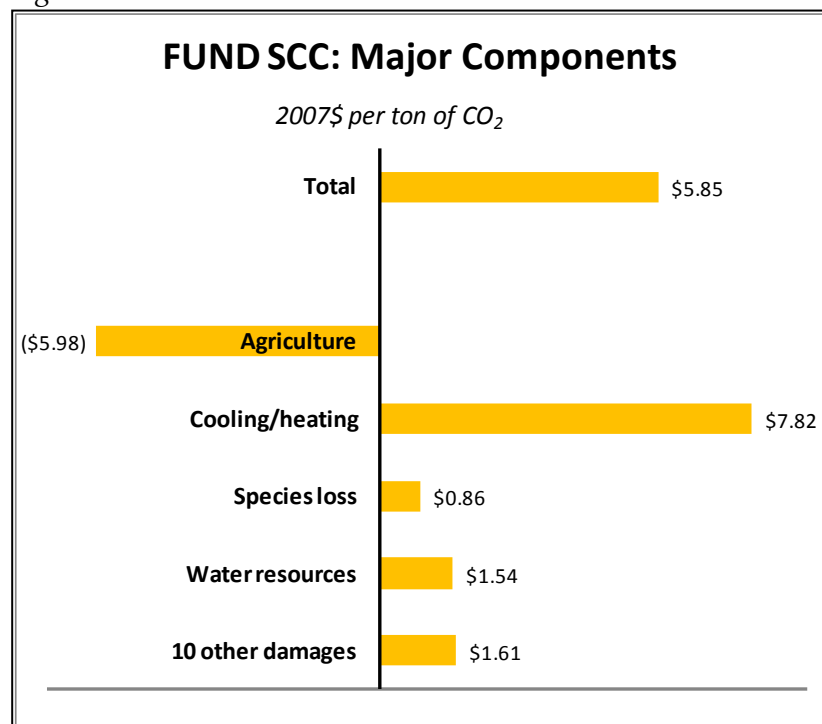
Disaggregating FUND damages

FUND presents separate calculations for 15 major impact categories (of which several, including health and agriculture, include separate calculations for multiple subcategories). Two of the major categories are closely related to each other, namely the increased costs for space cooling and decreased costs for space heating, as consequences of rising average temperatures. They are combined into a single cooling/heating category in the following presentation. The cooling/heating category is always a net cost of warming, since FUND's estimate of air conditioning costs increases with temperature more rapidly than its estimate of heating costs decreases.

---

[4] PAGE has a more complex treatment of damages than DICE, making it difficult to repeat the same experiment with the PAGE damage function.

The agriculture and cooling/heating categories are the only large components of the FUND SCC estimate; the other 12 are quite small. Figure 1 shows the impacts of the most important categories, when running FUND in the Monte Carlo mode used by the Interagency Working Group.
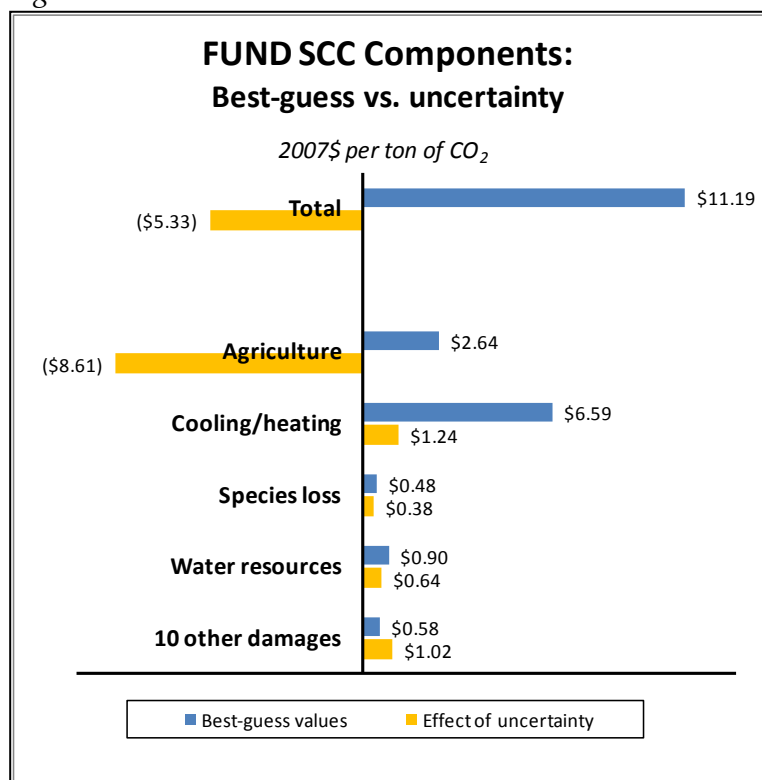
Figure 1

FUND SCC: Major Components

2007$ per ton of CO_2

Category	Value
Total	$5.85
Agriculture	($5.98)
Cooling/heating	$7.82
Species loss	$0.86
Water resources	$1.54
10 other damages	$1.61

FUND's $6 SCC estimate is the sum of a $6 net benefit in agriculture, a $8 net cost in cooling and heating, and a total of $4 of net costs in the other 12 damage categories combined. The largest of the other 12 are water resources and species loss; the remaining 10 categories, including sea-level rise, storm damages, wetland losses, human health, and migration impacts, amount to a combined total of less than $2 per ton of CO_2. One of the 10 smaller categories, forestry impacts, is a very small net benefit; the others are all small net costs.

Note that the impact of cooling and heating is greater than the SCC as a whole. Thus under the Working Group assumptions, FUND estimates that all impacts of climate change, excluding the increased costs of air conditioning, would amount to a net benefit to the world.

FUND also offers the option of calculation in "best-guess" mode, fixing all the Monte Carlo parameters at their modal values. Running the Working Group analysis in best-guess mode produces a SCC estimate of $11, compared with $6 in the Monte Carlo analysis. It seems natural to define the effect of uncertainty in FUND as the difference between the Monte Carlo values and the best-guess values. Using that definition, the estimates shown in Figure 1 can be broken down as follows:

Figure 2

The sum of the two bars for each category in Figure 2 is the value shown for that category in Figure 1. The effect of uncertainty is positive (increases the SCC) in all cases except agriculture. Uncertainty is only a small part of the impact of cooling and heating, and about half of the impact for the 12 smaller categories. In agriculture, however, the best-guess impact is a small positive amount, or net cost, while the effect of uncertainty is a larger negative, or net benefit.
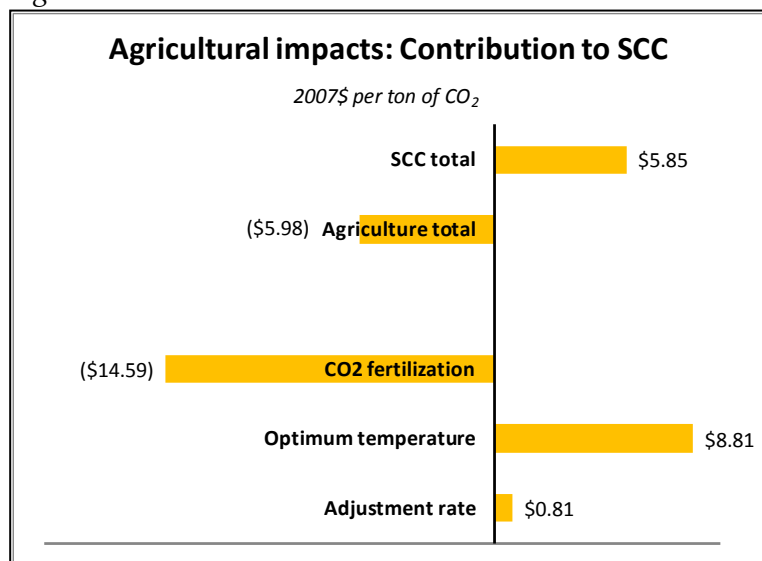
Agricultural impacts

In view of the importance of agricultural impacts, as seen in Figure 1 and 2, it is worth taking a closer look at this category. FUND models agricultural impacts as the sum of three effects:

- The CO_2 *fertilization effect* assumes that agricultural production is proportional to the logarithm of CO_2 concentrations. This is always a net benefit of climate change (i.e., reduction in the SCC).
- The *optimum temperature effect* assumes that agricultural production is a quadratic function of temperature, reaching a maximum at a temperature with a most likely value somewhat above current levels. The sign of this effect can vary.
- The *adjustment rate effect* assumes that agricultural production is decreased by adjustment costs, which are proportional to the rate of change in temperature; this is always a small net cost (increase in the SCC).

Using the same methodology, these effects can be turned off one at a time to determine their effects on the SCC. The results, corresponding to Figure 1, are shown in Figure 3, with the Working Group SCC

and the total agricultural impact repeated from Figure 1, for ease of comparison. The negative (beneficial) impact in agriculture is entirely due to CO_2 fertilization, which is estimated to provide a net benefit of more than $14 per ton of CO_2 emissions.

Figure 3

Agricultural impacts: Contribution to SCC

2007$ per ton of CO_2

SCC total — $5.85
($5.98) Agriculture total
($14.59) CO2 fertilization
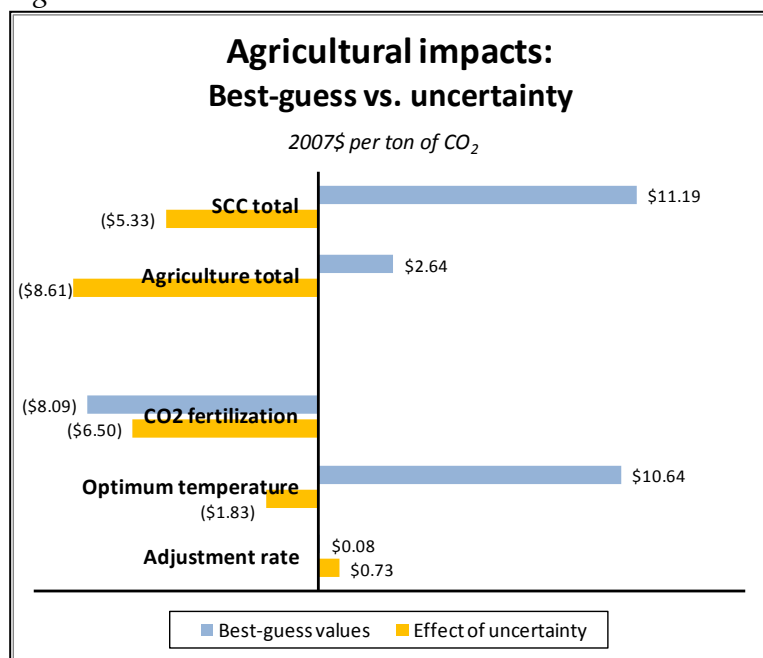Optimum temperature — $8.81
Adjustment rate — $0.81

In Figures 1 and 2, the components of the SCC add up precisely to the total; in Figure 3, the subcategories of agricultural impacts do not add up exactly to the total for agriculture, due to a constraint on these impacts in the FUND software.[5] The difference, however, is only about $1.

The best-guess values and the effects of uncertainty can be compared for the three agricultural subcategories, as was done for the broader categories in Figure 2. The results are presented in Figure 4. For CO_2 fertilization, both the best-guess value and the effect of uncertainty are net benefits (reductions in the SCC). The adjustment rate impact is essentially entirely a result of uncertainty; this is not surprising, since the most likely values for the adjustment rate parameter are very small, and are less than the standard deviation for every region. For the optimum temperature impact, the best-guess value and the effect of uncertainty have opposite signs – unlike the other agricultural subcategories, or the other impact categories shown in Figure 2. The best-guess optimum temperature impact is a net cost (increase in the SCC), while uncertainty about this impact reduces the SCC.

---

[5] FUND limits each region's total agricultural impacts to being no greater than the contribution of agriculture to the region's GDP. This constraint is not binding in the best-guess run, but it is in some of the Monte Carlo iterations. In the presence of this constraint, the impacts of the individual agricultural effects do not sum to the total agricultural impact. Thus the best-guess estimates for the three agricultural effects sum to the total agricultural best-guess value, but the same is not true of the Monte Carlo estimates.

Figure 4

Modeling agricultural impacts: Three issues

Further examination of FUND's agricultural calculations reveals three issues that need attention: two involving optimum temperature impacts, and one involving CO_2 fertilization impacts.

Risk of division by zero

The manner in which the optimum temperature effect is modeled in FUND 3.5 could cause division by zero for a plausible value of a Monte Carlo parameter. The equation for the optimum temperature impact, modeled as a percentage change in agricultural output, is (in slightly simplified notation):

$$(2) \qquad Impact = \frac{-2AT^{opt}}{10.24-6.4T^{opt}}T + \frac{A}{10.24-6.4T^{opt}}T^2$$

This is calculated for each time period and region. T is the average change in temperature, a global variable, and T^{opt} is the optimum temperature for agriculture. Both A and T^{opt} are Monte Carlo parameters, specified separately for each region.

In equation (2), the denominators of both fractions would be zero if T^{opt} =1.6. This is not a problem in FUND's best-guess mode; the regional values of T^{opt} are never equal to 1.6. The closest is 1.51, and most are much farther away. In Monte Carlo mode, however, T^{opt} is an unconstrained, normally distributed variable; the critical value of 1.6 is within 0.25 standard deviations of the mean for every region. This implies that it will be reasonably common to draw a value very close to 1.6, making the denominator very small and the impact very big. In such cases, the magnitude of the impact will

depend primarily on how close to 1.6 the value of T^{opt} turns out to be. Ironically, this problem could become more severe as the number of Monte Carlo iterations increases, since the likelihood of coming dangerously close to the critical value increases. (In the Working Group analysis, there are 10,000 iterations, each involving selection of 16 values of T^{opt}, one for each region.)

A fix for this bug is planned for the next version of FUND.[6] The anomaly is unfortunately present, however, in the versions that have been used in the past, including version 3.5, which was used for the Working Group's calculation of the SCC. In FUND versions 2.8 and 3.3, the earlier versions for which documentation is available on-line, the optimum temperature impact is defined by an equation with the same structure as (2), but with the denominators of the two fractions equal to $(1 - 2T^{opt})$. Thus the critical value that would cause a zero denominator was $T^{opt} = 0.5$ in FUND version 3.3 and earlier.

Two simple ways of removing the problem would imply similar changes in the FUND estimate of the SCC. (Both methods affect both uses of damage calculations in FUND, as described in the methodology section above; unlike other recalculations of the SCC, they are based on slightly different income trajectories from the Working Group's central estimate.) First, FUND can be run with T^{opt} fixed at its best-guess value for each region; that is, equation (2) is unchanged, but T^{opt} is no longer a Monte Carlo parameter.[7] Everything else about the model, including the definition of A in equation (2) as a Monte Carlo parameter, is also unchanged. This change has no effect on the best-guess value, but increases the Working Group's central estimate of the SCC by more than $10, from $5.85 to $16.21.

Alternatively, equation (2) can be modified to use the global average value of T^{opt}, roughly 1.28, in the denominator of both fractions. The denominator becomes equal to 2.056, so the equation becomes

$$(3) \quad Impact = \frac{-2AT^{opt}}{2.056}T + \frac{A}{2.056}T^2$$

In this variant, both A and T^{opt} are still Monte Carlo parameters, but T^{opt} no longer appears in the denominator. This change alone increases the best-guess value of the SCC only slightly, from $11.19 (as shown in Figure 2) to $11.68. The Monte Carlo estimate, corresponding to the Working Group's $5.85, becomes $17.98, or an increase of $12.

Thus two different ways of eliminating the problem in the optimum temperature equation, making no other changes, would raise the FUND estimate of the SCC by $10-$12.

Implausible temperature ranges

In addition to the potential problem of near-zero denominators, the optimum temperature equation employs an extremely wide range of variation in its Monte Carlo analysis. Table 1 presents data, from the FUND 3.5 data tables, on optimal temperatures for the model's 16 regions of the world: the means and standard deviations of the normal distributions used in the Monte Carlo analysis, and a calculation of the 95 percent confidence intervals (the mean plus or minus 1.96 standard deviations). The means

---

[6] Personal communication, David Anthoff, December 2010.
[7] We implemented this change (and a similar one described later in this article) by setting the standard deviations to zero in the FUND data file.

are smaller than the standard deviations in every case, much smaller in most cases; if this is the best information available about optimum temperatures, one could argue that they may not be significantly different from zero. (The same could be said, for the same reason, of the agricultural adjustment rate effect parameter; but as seen in Figures 3 and 4, the adjustment rate effect is close to zero in any case.)

Table 1

Optimal temperature for agriculture in FUND 3.5					
		(^oC above 1990)			
				95% confidence interval	
	Mean (μ)	Standard deviation (σ)		μ - 1.96σ	μ + 1.96σ
USA	1.09	4.14		-7.02	9.20
CAN	2.92	7.64		-12.05	17.89
WEU	0.79	3.29		-5.66	7.24
JPK	0.98	6.61		-11.98	13.94
ANZ	2.00	8.00		-13.68	17.68
EEU	1.31	2.73		-4.04	6.66
FSU	1.46	2.44		-3.32	6.24
MDE	1.32	2.03		-2.66	5.30
CAM	1.05	3.60		-6.01	8.11
SAM	0.35	8.82		-16.94	17.64
SAS	1.13	2.41		-3.59	5.85
SEA	0.70	5.12		-9.34	10.74
CHI	1.43	2.49		-3.45	6.31
NAF	1.20	2.74		-4.17	6.57
SSA	1.22	2.76		-4.19	6.63
SIS	1.51	2.92		-4.21	7.23

The width of the confidence intervals in Table 1 appears to exceed physically plausible temperature ranges for agriculture. FUND asserts 95 percent confidence that the optimal temperature for agriculture in South America, for instance, is between 17^{o}C below and almost 18^{o}C above 1990 levels. For the United States, the corresponding range is from -7^{o}C to +9^{o}C. The upper end of the 95 percent confidence interval is more than 5^{o}C above 1990 temperatures everywhere, and more than 10^{o}C above 1990 in five regions.

Monte Carlo analysis across these intervals – including the even higher "optimum temperatures" that will be chosen for each region in 250 of the 10,000 iterations – would seem to be exploring hypotheses about the state of the world that could safely be ruled out in advance. In each Monte Carlo iteration that selects a very high optimum temperature, FUND calculates a double benefit from climate change: both the fertilization from increasing CO_2 concentrations, and the increasing (but still sub-optimal) temperature, are estimated to have separate, positive effects on agriculture. Since FUND has a lower bound on agricultural damages (see footnote 5), but no upper bound on agricultural benefits, Monte

Carlo analysis across an excessively wide range of possibilities increases the reported average agricultural benefits.

Truncated distributions for carbon fertilization

A different problem arises in FUND's Monte Carlo analysis of carbon fertilization. FUND models this effect using the following equation (again using simplified notation):

$$(4) \qquad Q = \frac{\gamma}{\ln(2)} \ln\left(\frac{C}{275}\right)$$

Here Q is the percentage increase in agricultural output due to carbon fertilization, C is the atmospheric concentration of CO_2 in ppm, and γ is the carbon fertilization effect at 550 ppm.[8] The strength of the effect, i.e. γ, is a Monte Carlo parameter specified separately for each region. The probability distribution used for this parameter is a "truncated normal" distribution – the portion of the normal distribution for which the parameter value is positive.[9]

FUND makes extensive use of this technique; the final section of the data tables for FUND 3.5 shows that 51 of the model's 73 Monte Carlo parameters are assumed to have truncated normal distributions, constrained to always use positive, or in some cases always negative, values. This makes little difference in practice when the mean is much larger than the standard deviation. If the mean exceeds two standard deviations (in absolute value), then truncating the normal distribution at zero eliminates less than 2.5 percent of the area under the curve; in such cases the mean of the truncated distribution is nearly identical to the reported value for the unconstrained normal curve.

When the mean is not large relative to the standard deviation, however, truncating at zero produces a distinctly non-normal distribution, with a different mean than the unconstrained curve. This occurs in a number of places in FUND, but is never discussed in the model's documentation.

In the case of the carbon fertilization parameter, γ, the mean is greater than two standard deviations in seven regions of the world, and between one and two standard deviations in another five. For the remaining four regions, the United States, Canada, Japan/Korea, and Southeast Asia, the mean is less than the standard deviation. In these cases, as shown in Table 2, the true mean of the truncated normal distribution is more than 50 percent higher than the reported mean of the unconstrained curve.[10]

---

[8] The factor of ln(2) in the denominator, which allows this simple interpretation of γ, is present in the FUND software but omitted in the technical description of version 3.5.

[9] The implementation of this feature in the FUND software consists of drawing a value from the unconstrained normal distribution; if it is in the truncated region (a negative value, in this case), another value is drawn, repeating the process until an acceptable value is obtained.

[10] The mean of the truncated distribution was calculated using Excel statistical functions: using μ for the mean and σ for the standard deviation of the unconstrained normal distribution, the area removed by truncation is *TRUNC* = NORMDIST(0,μ,σ,TRUE); the mean of the truncated distribution is NORMINV((1+*TRUNC*)/2,μ,σ). See Figure 5 for an example.

Table 2

Carbon fertilization parameter in FUND 3.5				
(fractional increase in output at 550 ppm)				
	Unconstrained normal distribution			*Truncated distribution*
Region	**Mean**	**Standard deviation**		**True mean**
USA	0.089	0.148		0.141
CAN	0.040	0.065		0.062
WEU	0.154	0.118		0.168
JPK	0.232	0.366		0.355
ANZ	0.105	0.085		0.116
EEU	0.095	0.051		0.097
FSU	0.067	0.055		0.075
MDE	0.094	0.027		0.094
CAM	0.164	0.054		0.164
LAM	0.060	0.050		0.067
SAS	0.058	0.016		0.058
SEA	0.085	0.418		0.316
CHI	0.192	0.061		0.192
MAF	0.073	0.019		0.073
SSA	0.051	0.022		0.051
SIS	0.238	0.086		0.238

Figure 5 contrasts the unconstrained and truncated distributions, and their respective means, for Southeast Asia. The dashed curve shows the negative portion of the unconstrained normal curve, generated with FUND's reported mean and standard deviation. The solid red line is the truncated normal distribution; its mean – that is, the mean of the distribution which FUND actually uses – is almost four times as large as the reported value.

11

Figure 5

Mean of full distribution μ_F = 0.085
Mean of truncated distribution μ_T = 0.316
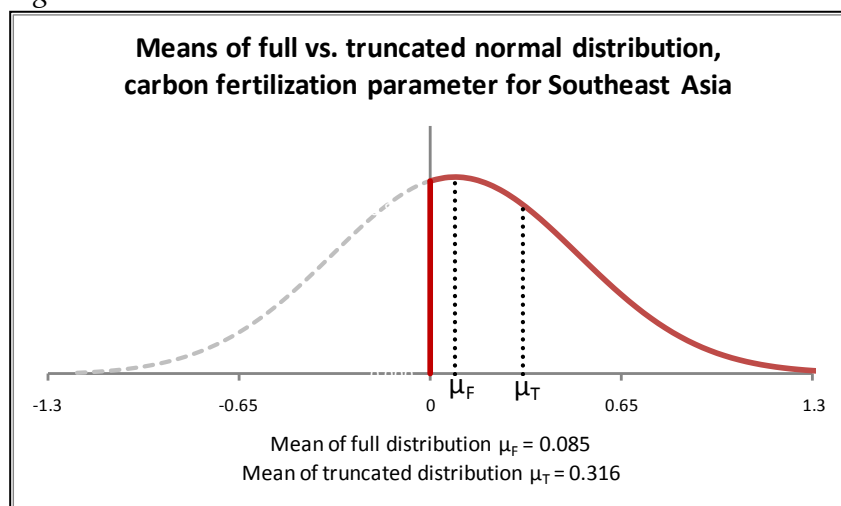
When the truncation causes a significant change in the shape of the curve, as in Figure 5, there is no obvious theoretical rationale for the use of the truncated normal distribution. A rough attempt to measure the impact of this problem can be made by constraining FUND to use its best-guess values for γ – that is, the values shown in the first column of Table 2. With this change, the SCC becomes $21.60, an increase of almost $16 above the Working Group value. (Like the revised SCC estimates presented in section 3.3, this estimate changes the baseline income trajectory, and therefore is not directly comparable to the estimates in Figures 1-4.)

Implications: the need for updated estimates

Since the FUND model remains important in the ongoing discussion of climate policy, there is a need to update and improve its damage estimates. In the area of agricultural impacts, FUND's technical description states that the model's estimates are calibrated to research results published in 1992-1996. There has been a substantial advance in the understanding of agriculture and climate change since 1996, which might lead to different estimates.

Early studies of carbon fertilization, usually done in greenhouses, suggested that it would lead to very large gains in agricultural yields. Recently, however, more realistic outdoor experiments have suggested that the benefits will be much smaller, perhaps half the size of the earlier estimates (Long et al. 2006; Leakey et al. 2009). A recent economic analysis of agriculture and climate change concluded that an increase in atmospheric concentration to 550 ppm of CO_2 would, on average, increase agricultural yields by 9 percent (Cline 2007).

When a simple carbon fertilization relationship is assumed to apply to all future CO_2 concentrations, there is a risk of out-of-sample forecasting: as concentrations rise, in high-emission climate scenarios, do yields keep rising forever? An unbounded logarithmic relationship between CO_2 concentrations and yields, as assumed in FUND, means that each doubling of CO_2 concentrations produces the same increase in agricultural output. Yet there is very little empirical information available about yields at higher concentrations.

A more cautious modeling approach might assume moderate yield gains, along the lines of Cline (2007), for the initial increases in CO_2 concentration, but little or no further gains thereafter. This would reduce the large net benefits which FUND currently estimates from CO_2 fertilization, particularly in high emission, business-as-usual scenarios.

The optimum temperature effect, as modeled in FUND, makes agricultural output a quadratic function of temperature (see equation (2) above); even with the simplest fixes for the division-by-zero problem, as proposed in the last section, the relationship is still quadratic. This implies perfect symmetry between the impacts of higher- and lower-than-ideal temperatures: with a quadratic relationship, the projected yield is necessarily the same at 1^o above and 1^o below the optimum. Again, recent research suggests a different pattern.

In a detailed empirical study of the effects of temperature on U.S. corn, soybeans, and cotton yields, Schlenker and Roberts (2009) found very slow, small increases in yields on the way up to the optimum temperature (which was 29^oC for corn, 30^oC for soybeans, and 32^oC for cotton), followed by rapid declines in yields above the optimum. For corn, replacing 24 hours of the growing season at 29^oC with 24 hours at 40^oC causes a predicted yield decline of about 7 percent.

Their results do not at all resemble a quadratic relationship; a closer approximation would be a horizontal line (constant yield) up to the optimum temperature, followed by a steep drop-off in yield at higher temperatures. This would require a different functional form for the optimum temperature effect, in place of equation (2). Schlenker and Roberts find no evidence of successful adaptation, such as development of heat-resistant crop varieties, in parts of the country which have long been above the optimum temperatures for much of the growing season.

Corn, soybeans, and cotton are three of the world's highest-value crops, and the United States produces a significant fraction of global supply, including 41 percent of corn and 38 percent of soybeans (Schlenker and Roberts 2009). Thus this is not just a case study, but a description of a large part of world agricultural production. Use of the Schlenker and Roberts curves, in place of FUND's current quadratic relationship between yield and temperature, would have a major effect on the estimates of agricultural impacts of climate change: it would reduce the large estimated gains from warming, particularly in the Monte Carlo iterations where FUND currently picks very high optimum temperatures.

Conclusions

Model results matter. The estimate of the SCC adopted by U.S. government agencies, for use in calculations such as cost-benefit analyses of proposed regulations, is based on the results of three models of climate economics – of which FUND is the most complex and least understood. Models that play such a prominent role need to be transparent, widely understood, and up to date and consistent with the latest empirical research.

This paper has introduced a software innovation that increases the transparency of the FUND model: switches that allow individual damage categories to be turned on and off, in order to understand their relative contributions to the final results. FUND's $6 SCC estimate, lower than some other models, is the sum of an estimated net benefit in agriculture, a net cost in heating and cooling, and very small net costs in all other areas.

All of these areas should be examined and updated as needed; some widely discussed climate impacts, such as sea-level rise and extreme weather events, are surprisingly estimated to add almost nothing to the SCC in FUND. This could be an important, counterintuitive result about the magnitude of the empirical evidence, or it could be an indication that FUND's impact estimates are in need of revision.

In the area of agriculture, FUND currently relies on research from 1996 or earlier to estimate a large net benefit from CO_2 fertilization, an optimum temperature effect on yields, and a small effect from the rate of temperature change. The first two, which account for virtually the entire agricultural estimate, are both in need of revision. Newer research suggests smaller benefits from CO_2 fertilization, and says nothing about whether these benefits continue at very high concentrations. A flaw in FUND's optimum temperature equation needs to be fixed, to prevent the risk of division by zero; and the quadratic shape of that equation is inconsistent with recent research on temperature and yields.

Since model results matter, so do the damage calculations used inside the models. If FUND used the DICE damage function, it would roughly agree with DICE about the SCC, estimating it at $31 instead of $6 under the assumptions made by the Interagency Working Group. The quick fixes to the individual modeling problems identified in sections 3.3 and 3.4 above would raise the FUND estimate of the SCC by $10-$16, which represents half or more of the difference between the FUND and DICE estimates.

This does not mean that the DICE damage function, or any of the modifications to the FUND equations discussed here, would produce the right estimate of the SCC. The authors of FUND have, quite reasonably, responded that simply turning off the Monte Carlo variation on one of their parameters is not an appropriate way to revise the model.[11] Nonetheless, the problems identified here require attention. Much more careful work, including examination of damage categories beyond agriculture, should be done to produce an adequate revision of FUND. The point is simply that problems in model specification and methodology, and failure to update the empirical evidence used in the model, can have relatively large effects on the results. Since model estimates are being treated as establishing a precise SCC value, suitable for use in policy analysis, it is essential to revisit and revise those estimates, and the assumptions and inputs behind them, on a regular basis.

---

[11] Personal communication, David Anthoff, December 2010.

References

Cline, W.R. (2007). *Global Warming and Agriculture: Impact Estimates by Country*. Washington, DC: Center for Global Development & Peterson Institute for International Economics. Available at http://www.cgdev.org/content/publications/detail/14090.

Interagency Working Group on Social Cost of Carbon (2010). "Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866." In *Final Rule Technical Support Document (TSD): Energy Efficiency Program for Commercial and Industrial Equipment: Small Electric Motors*. Washington, DC: U.S. Department of Energy, p. Appendix 15A. Available at http://www1.eere.energy.gov/buildings/appliance_standards/commercial/pdfs/ smallmotors_tsd/sem_finalrule_appendix15a.pdf.

Leakey, A. D. B., Ainsworth, E. A., Bernacchi, C.J., et al. (2009). "Elevated CO2 effects on plant carbon, nitrogen, and water relations: six important lessons from FACE." *Journal of Experimental Botany* 60(10), 2859-2876. DOI: 10.1093/jxb/erp096.

Long, Stephen P., Ainsworth, Elizabeth A., Leakey, Andrew D.B., et al. (2006). "Food for Thought: Lower-Than-Expected Crop Yield Stimulation with Rising CO2 Concentrations." *Science* 312(5782), 1918-1921.

Nordhaus, W.D. (2008). *A Question of Balance: Economic Modeling of Global Warming*. New Haven, CT: Yale University Press.

Schlenker, W. and Roberts, M.J. (2009). "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change." *Proceedings of the National Academy of Sciences* 106(37), 15594 -15598. DOI: 10.1073/pnas.0906865106.