CESifo

Working Papers

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

Thijs Dekker
Rob Dellink
Janina Ketterer

CESifo
Center for Economic Studies & Ifo Institute

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

Abstract

International climate negotiations take place in a setting where uncertainties regarding the impacts of climate change are very large. In this paper, we examine the influence of increasing the probability and impact of large climate change damages, also known as the 'fat tail', on the formation of an international mitigation agreement. We systematically vary the shape and location of the distribution of climate change damages using the stochastic version of the applied game-theoretical STACO model. Our aim is to identify how changes to the distributional form affect the stability of coalitions and their performance. We find that fatter upper tails increase the likelihood that more ambitious coalitions are stable as well as the performance of these stable coalitions. Fatter tails thus imply more successful, or 'fatter', international climate agreements.

JEL-Code: C790, H870, Q540.

Thijs Dekker
Institute for Environmental Studies (IVM)
Leeds University / Leeds LS2 9JT / UK
t.dekker@leeds.ac.uk

Rob Dellink
Wageningen University
6700 EW Wageningen / The Netherlands
rob.dellink@wur.nl

Janina Ketterer
European Bank for Reconstruction and Development
EC2A 2JN London / UK
ketterej@ebrd.com

December 2012

1. Introduction

Avoiding dangerous interference with the climate system requires substantial reductions in global emission levels (IPCC, 2007). Large regional differences in the costs and benefits from emission reductions complicate international coordination of these mitigation efforts. The efficient level of mitigation efforts can be attained through global cooperation, i.e. by forming an international environmental agreement (IEA) in which all countries participate. Due to the public good nature of the climate system, however, countries benefit from the mitigation efforts of other countries even if they do not contribute by reducing their own emissions. Hence, free-riding behaviour becomes attractive and motivates some countries to stay outside the coalition. As there is no supra-national authority to decide *which* countries should undertake mitigation efforts and *how much* emission reductions should be conducted, international negotiations have to rely on attaining a self-enforcing agreement.

This debate's complexity is further amplified by the inherent uncertainties related to the climate system. Uncertainties and risks in the climate system are compounded with uncertainties surrounding the economic evaluation of the impacts of climate change.[4] One key uncertainty in evaluating climate policies relates to the (small) possibility of very large damages. The more likely these events are, the more they influence the optimal policy response. The sensitivity of policy responses outcomes to high-consequence, low-probability impacts has been discussed in various climate models (e.g. Dietz, 2009; Anthoff and Tol, 2010).

In this paper we augment this line of research by examining the influence of increasing the probability of large climate change damages, also known as the 'fat tail', on the stability and performance of an IEA.[5] What will happen to the size of IEAs if the likelihood of very high

---

[4] Uncertainties and risk are inherent in the climate system. In the STACO model we track the influence of risk and uncertainty by performing Monte Carlo simulations, specifying different distributional forms and parameters. We do not make a strict distinction between the terms uncertainty and risk, but use both terms to refer to an unknown impact of climate change.

[5] Weitzman (2009a, 2009b) initiated the discussion on 'fat tails' by criticising Integrated Assessment models (IAMs) for underestimating climate change damages. With reference to Nordhaus (2009), we assume explicitly that IAMs remain a valid tool as Weitzman's invalidating conditions do not apply to the wide range of climate scenarios investigated here, as long as the model allows for some mitigation action and carefully specifies the distributional forms to represent uncertainty. Within these boundaries we apply alternative scenarios and examine their impact on the stability and performance of an IEA. The expression 'fat tail' is used to describe a distribution in which high impacts have a higher probability than can be expected based on a normal distribution.

climate damages increases? Are countries more likely to join an international coalition when they face greater uncertainty and risk about climate change damages? Can high, but unlikely, impacts ensure larger (joint) mitigation efforts? Using the stochastic version of the STACO model (Dellink et al., 2008), we analyse the extent to which IEAs are affected by varying degrees of climate change risks.

The basic structure of the STACO model consists of interacting regions that (i) choose to join an international mitigation agreement or not, and (ii) choose their optimal mitigation policy given the coalition formed. Regions are characterised by their abatement costs and damage cost functions and linked via global abatement (mitigation) efforts. Dellink et al. (2008) introduce uncertainty in the STACO model by specifying a set of stochastic parameters in the cost and benefit functions. Moreover, they introduce the concept of stability likelihood (SL), linking uncertainties about costs and benefits of mitigation efforts to the stability of IEAs. In a set of basic sensitivity analyses, Dellink et al. (2008) show that the value of the SL for a specific coalition mainly depends on two elements: the variance in regional benefit shares of mitigation as well as the variance and level of the global benefit parameter.[6]

Dellink and Finus (2012) extend the analysis of Dellink et al. (2008) by introducing transfers among coalition members. They focus their analysis on the impacts of learning effects, i.e. how the resolution of uncertainty affects stability, building upon Na and Shin (1998); Kolstad (2007); Kolstad and Ulph (2008); Kolstad and Ulph (2011). Dellink and Finus (2012) show that the common conclusion from the papers by Kolstad and Ulph that learning leads to worse outcomes in terms of welfare and environmental quality does not have to hold when (optimal) transfers are available.

Compared to Dellink et al. (2008) and Dellink and Finus (2012), this paper shifts focus from general uncertainty analysis to assessing the impact of the risk of very high climate damages. It provides a more rigorous analysis of the impact of the probability of highly disruptive climate change damages on the stability and performance of climate coalitions. Fat tails in the distribution of mitigation benefits are introduced by (i) increasing the variance of the global benefit parameter, and (ii) opting for asymmetric distributional shapes. By selecting the

---

[6] In the STACO model benefits are characterised by a stream of prevented climate damages due to mitigation efforts.

appropriate parameter values we are able to increase the density in a specific tail. Based on these systematic variations in the properties of the global benefit parameter, we subsequently conduct a regression analysis. Our aim is to assess the impact of fatter upper tails of the distribution on the SL and related performance measures (as defined in Section 2) of all coalitions in the STACO model. We find that many coalitions are only stable in the trivial case when climate change has a net positive impact on a global scale, because joining the coalition does not require any mitigation efforts in this case. We find that by shifting mass to the upper tail, i.e. by increasing the likelihood of very high damages, both strict stability likelihood and performance of an international mitigation agreement increases. In other words, fatter tails do imply more successful, or 'fatter', international climate agreements.

The paper is structured as follows. Sections 2 and 3 give an overview of the theory on coalition formation and the properties of the applied STACO model. Section 4 outlines the scenarios used in this study, focusing on distributional changes of the global benefit parameter and the introduction of fat tails. Thereafter, Section 5 presents the results of an econometric analysis quantifying the impact of increased climate change risks on the stability likelihood and the performance of coalitions. Implications and limitations are discussed in concluding Section 6.

2. Coalition formation

2.1. A model framework for coalition formation

We adopt a standard framework of cartel formation with single deviations, following Barrett (1994). Consider a set of N heterogeneous players, each representing a country or world region. Each player is involved in a two-stage coalition formation game. In the first stage, players decide whether to become a member of an IEA or to remain an outsider. Announcement $c_i = 1$ implies that player i joins the coalition, while $c_i = 0$ indicates player i stays out of the coalition. The coalition structure can be summarised by the announcement vector $c = (c_i \ldots, c_N)$. Players announcing 1 are called coalition members and together they form coalition $k = \{i \mid c_i = 1, \forall i = 1,\ldots,N\}$. The terms coalition structure c and coalition k can be used interchangeably. We denote the set of possible coalitions by K.[7] In the second stage,

---

[7] Note that there are 2^N announcement vectors, but there are only 2^N-N different coalitions as coalitions of only one member are trivial.

players decide on their mitigation (abatement) levels as illustrated by the abatement vector $q = (q_1,...,q_N)$. Due to the public good nature of climate change, the benefits from mitigation (i.e. prevented climate damages) are driven by global abatement levels, while costs of mitigation are borne locally. Thus, each player's payoff $\pi_i(q,b)$ depends on the global level of abatement, its own abatement level and a vector of model parameters b.[8]

The game is solved backwards assuming that strategies in each stage form a Nash equilibrium. For the second stage, this entails that a Partial Agreement Nash equilibrium is formed between coalition members k and the non-signatory players J (Chander and Tulkens, 1995). The coalition acts as one player and maximises its joint payoff. Simultaneously, each non-signatory player optimises its own payoff by selecting its abatement level conditional on the abatement levels of coalition k and other non-signatory players $-J$.

$$\sum_{i \in k} \pi_i(q_k^*, q_J^*, b) \geq \sum_{i \in k} \pi_i(q_k, q_J^*, b) \forall i \in k \tag{1}$$

$$\pi_i(q_k^*, q_i^*, q_{-J}^*, b) \geq \pi_i(q_k^*, q_i, q_{-J}^*, b) \forall i \in J \tag{2}$$

Equations (1) and (2) describe the necessary conditions for a Nash equilibrium in stage two, where q_k is the abatement vector of coalition k, q_J the abatement vector of all non-signatories; q_i is the abatement of player i if it is a non-signatory, and q_{-J} the abatement vector of all other fringe players except i. An asterisk denotes equilibrium strategies. This implies that the behaviour of non-signatories towards all other regions is selfish and non-cooperative; signatories behave in a cooperative way towards their fellow coalition members, but non-cooperative towards outsiders. The equilibrium strategy vector q^* corresponds to the classical ´social or global optimum` if coalition k comprises all countries, i.e. the grand coalition, and corresponds to the classical Nash equilibrium if coalition k comprises only one member or is empty. Any inefficiency stems from the fact that k is not the grand coalition. Olieman et al. (2006) prove that, in the context of our simulation model, the equilibrium abatement strategy vector q^* is unique for every coalition and a given vector of parameters. Consequently, there is a unique equilibrium payoff for every coalition $v_i(k,b) = \pi_i(q^*(k,b))$.

---

[8] Note that the payoff function is purely a monetary measure and not a fully specific utility function. A non-linear utility function would affect the preferences of countries to join the coalition, but is beyond the scope of the current paper.

Following d'Aspremont et al. (1983), the first stage Nash equilibrium requires that no signatory that announced $c_i = 1$ should have an incentive to change its announcement to $c_i = 0$ and no non-signatory that announced $c_i = 0$ should want to announce $c_i = 1$, given the announcement of all other players. The former is labelled internal stability and the latter external stability. The stability condition can be summarised compactly by $f(k,b)$, which assigns the value 1 to a stable coalition (i.e. stable announcement vector) and the value 0 to an unstable coalition. In Equation (3), k_{-i} represents the coalition where region i changed its announcement vector. It is worth noting that for any given set of parameters b, this function allows for either a unique stable coalition, multiple stable coalitions or no stable coalition at all.[9]

$$f(k,b) = \begin{cases} 1 & v_i(k,b) - v_i(k_{-i},b) \geq 0 \ \forall i \in N \\ 0 & otherwise \end{cases} \quad (3)$$

We compute the average aggregate valuation, i.e. global payoff, over the (non-dominated) coalitions by (4), where S is the number of non-dominated stable coalitions:

$$\overline{v}(b) = \frac{\sum_{k \in K} f(k,b) \sum_{i=1}^{N} v_i(k,b)}{S} \quad (4)$$

The formation of large stable coalitions may be hampered in the absence of transfers (Carraro and Sinisalco, 1993). Due to the asymmetry across players, some benefit more than others from a specific coalition. Signatory players with high marginal benefits have an incentive to share a part of their surplus with other signatory countries to make the coalition stable. In other words, a transfer scheme may exist that makes coalition k internally stable. The type of transfer scheme plays a large role in determining the stability of coalitions (Carraro et al., 2006; Nagashima et al., 2009). In this paper, we follow Carraro et al. (2006) and Weikard

---

[9] We check the stability condition by changing the announcement vector of one player at a time. Due to these single deviations, multiple stable equilibria are possible. In the case of multiple stable coalitions, we assume that each of them is equally likely to occur. There is a probability that one of them is (Pareto-)dominated by another stable coalition. The STACO model controls for this by assigning the set of Pareto-dominated coalitions a zero probability of occurring. The remaining set of stable coalitions is used to evaluate the success of coalition formation.

(2009) by applying the concept of "optimal transfers" or, as labelled by Carraro et al. (2006), "an almost ideal transfer scheme". Under this transfer scheme every coalition member receives its free-rider payoff when unilaterally leaving the coalition, plus an (arbitrary) share of the surplus. The latter is the aggregate payoff of the coalition minus the sum of free-rider payoffs. Accordingly, transfers are only paid among coalition members, and these transfers balance.[10]

2.2. Uncertainty, the concept of stability likelihood, and performance indicators

In a deterministic model, the vector of the model parameters b may be based on empirical estimates and detailed bottom-up impact assessments. When contradictory information is available for a specific model parameter, a meta-analysis of empirical studies and other available information can be used to infer an appropriate distributional form to characterise the uncertainty about its value. We incorporate parameter uncertainty in the model by declaring a set of stochastic parameters. Each stochastic parameter b is replaced by an independent distribution $g(b|\theta)$, where the set of hyper parameters θ characterises the shape and support of the probability density function $g(\cdot)$.

The way players respond to uncertainty in the game depends on the extent to which they learn about the true parameter values. Kolstad et al. (2007) distinguish three types of learning. When 'no learning' effects occur, decisions in both stages are based on maximising expected values.[11] In the case of 'partial learning' players learn about the true parameter values after the first stage, but before declaring their abatement levels. Decisions in stage one are therefore still based on expected payoffs from each coalition in stage two. 'Full learning' implies that players learn about the true parameter values at the start of the game.

Each coalition is either stable or unstable under no and partial learning, because decisions in the first stage are based on expected values of stage two outcomes. Under full learning the stability of a coalition is conditional on the true parameter value(s), since all uncertainty is resolved before stage one. Conditional on the specific parameter values, the game reduces to a deterministic model under full learning. Ex ante the true parameter values are, however,

---

[10] The interested reader is referred to Weikard (2009) for more details on the sharing scheme. Note that when coalitional payoff is insufficient to compensate all free-rider payoffs, the coalition becomes internally unstable.
[11] The STACO model described in the next section is linear in parameters (but not in abatement levels). Hence, the expected payoff is equal to the payoff based on the expected parameter vector (Dellink et al., 2008).

unknown. The concept of stability likelihood (SL) applies to this situation. The SL can be interpreted as the probability that coalition k is stable:

$$SL_k = P(f(k,b)=1) = \int_\theta f(k,b) \cdot g(b \mid \theta) d\theta$$

We approximate the SL using Monte Carlo simulations, because the heterogeneity of players' payoff functions in the STACO model precludes the existence of an analytical solution to this integral.[12] By generating M samples from $g(.)$ we can approximate the SL for coalition k by

$$\hat{SL}_k = \frac{1}{M} \sum_{m=1}^{M} f(k,b^m) \,,$$

where b^m represents the draw for b in sample m. Accordingly, the value for the binary stability condition $f(k,b)$ may vary over the draws. A more detailed discussion of the SL concept and computation can be found in Olieman et al. (2006).

The no learning case provides a useful benchmark for comparing model results, since the model effectively reduces to a deterministic setting using mean values. Also under partial learning the fatness of the upper tail of climate damages increases the expected value of the global benefit parameter and associated abatement efforts, but limited information is obtained on how the shape and support of the distribution for the global benefit parameter affect the stability and performance of coalitions. The full learning case is more suitable for our research question, as the Monte Carlo analysis provides information on stability and performance at all realisations of the (set of) uncertain parameter(s).

Apart from the SL measure, average payoff values and average abatement efforts (again averaged over all M samples), we introduce two additional performance indicators for each coalition. Success of a coalition is defined as the product of its stability likelihood and the average global payoff over all M samples:

---

[12] We are not aware of any paper that provides analytical solutions of stable coalitions in the context of heterogeneous players even in the absence of uncertainty.

$$Success_k = S\hat{L}_k \cdot \frac{\sum_{m=1}^{M}\sum_{i=1}^{N} v_i\left(k, b^m\right)}{M} \tag{5}$$

For example, the grand coalition generates the highest payoff from mitigation efforts, but is likely to have a low SL due to free-riding incentives that increase with coalition size. Its success will therefore be limited. Other coalitions may be more successful due to a higher SL despite the reduction in payoff.

Our second indicator of performance, efficiency, corrects for a potential level effect related to payoffs. We relate the payoff of coalition k to the payoff of the grand coalition (GC), which is fully efficient, and the payoff of the all singletons coalition (AS) which is least efficient. Conditional on the SL, the closer a coalition's payoff is to the GC, the larger its efficiency.

$$Efficiency_k = S\hat{L}_k \cdot \frac{1}{M}\left(\frac{\sum_{m=1}^{M}\sum_{i=1}^{N} v_i\left(k, b^m\right) - \sum_{m=1}^{M}\sum_{i=1}^{N} v_i\left(AS, b^m\right)}{\sum_{m=1}^{M}\sum_{i=1}^{N} v_i\left(GC, b^m\right) - \sum_{m=1}^{M}\sum_{i=1}^{N} v_i\left(AS, b^m\right)}\right) \tag{6}$$

3. The STACO model

3.1. Calibration of the deterministic STACO model

In this section, the calibration of the applied model, called STAbility of COalitions (STACO) is described; see Nagashima et al. (2009) and Dellink and Finus (2012) for more details. The model comprises benefit and cost functions of abatement efforts of twelve world regions: USA (USA), Japan (JPN), European Union (EEC), other OECD countries (OOE), Eastern European countries EET), former Soviet Union (FSU), energy exporting countries (EEX), China (CHN), India (IND), dynamic Asian economies (DAE), Brazil (BRA) and 'rest of the world' (ROW). Region i's payoff from abatement is given by:

$$\pi_i = \sum_{t=1}^{T}(1+r)^{-t} \cdot \left(B_{it}(q_t) - C_{it}(q_{it}) + F_{it}\right) \tag{7}$$

where T denotes the time horizon $t = 1,...,T$; r the discount rate; B_{it} benefits from global abatement $q_t = \sum_{i=1}^{N} q_{it}$ and C_{it} abatement costs from regional abatement of CO_2 q_{it}. Receipts from international transfers (endogenously determined in the model as explained in Section 2.1) are denoted by F_{it}. The payoff function is expressed in terms of the net present value of the regional abatement levels over the time period of interest. In this case, we apply a period of 100 years, starting in 2011, and a constant discount rate r of 2%, which roughly resembles a zero pure rate of time preference. Following Nagashima et al. (2009), the benefit function is a linear approximation of a three-layer carbon cycle linking current global abatement activities to a stream of future avoided damages as proposed by Nordhaus and Zhang (1996).[13]

Benefits from abatement in period t, as defined in (8), equal the net present value (in period t) of future avoided damages.

$$B_{it}(q_t) = \sum_{z=t}^{T}(1+r)^{t-z} \cdot \left(D_{iz}(q_t = 0) - D_{iz}(q_t)\right) \tag{8}$$

Future avoided damages are calculated by contrasting the damage function D_{iz} under no abatement efforts in period t and under global abatement efforts q_t. The damage function links abatement efforts and climate impacts by assuming an exogenous path of radiative forcing from non-CO_2 greenhouse gases. It is defined by: $D_{iz}(q_t) = c_i + \omega_{z-t} \cdot q_t \cdot s_i \cdot \gamma \cdot Y_z$, where c_i is a scaling parameter that has no effect on benefits in (8) as it cancels out; $\omega_{z-t} \cdot q_t$ reflects the impact of abatement in period t on atmospheric concentrations in period z; s_i are (stochastic) regional damage shares and $\gamma \cdot Y_z$ is the (stochastic) scale parameter of global damages multiplied by (future period) global GDP.[14] In the following, we label γ the global benefit parameter as benefits arise from prevented climate damages. The benefit function is linear, as it combines a (more or less) quadratic relation between damages and temperature change with a (by approximation) log-linear relationship between the stock of greenhouse gases and temperature change, features which are not uncommon in IAMs (e.g. Nordhaus and Yang, 1996). This delivers a linear function for the relevant domain of the stock of greenhouse gases. The model thus captures the stock pollutant aspect with inertia in the climate system,

---

[13] While this may ignore much of the interactions that take place in the climate system, it suffices for our goal of valuing the benefits of abatement activities.

[14] Projections for GDP are taken from the MIT-EPPA model (Paltsev et al., 2005).

see Dellink et al. (2010) for more details. The impact of non-linear benefits has been explored using a deterministic version of the STACO model in Weikard et al. (2006), who find only minor implications for stability of coalitions.

Abatement costs C_{it} are formulated following Ellerman and Decaux (1998), adjusted for an exogenous technological progress parameter ($\varsigma = 0.005$) to reflect the dynamic nature of our model, with additional (stochastic) parameters α_i and β_i:

$$C_{it}(q_{it}) = \tfrac{1}{3} \cdot \alpha_i \cdot (1-\varsigma)^t \cdot q_{it}^3 + \tfrac{1}{2} \cdot \beta_i \cdot (1-\varsigma)^t \cdot q_{it}^2 \qquad (9)$$

The benefit and cost functions are fully specified to obtain a payoff function that can be used for the numerical analysis, since analytical results cannot be obtained for the STACO model (Finus et al., 2006). The chosen functional forms are relatively flexible and commonly used functional forms. One essential property of these functions is that the resulting payoff function is concave, i.e. there are diminishing returns to abatement.

3.2. Uncertainty in the STACO model and strict stability likelihood

In this paper we take uncertainty about the climate and economic systems into account by specifying a distribution for the key parameters that describe costs and benefits of abatement action: [γ, α, β, s].[15] For the region specific parameters α_i, β_i, and s_i we use exactly the same set-up as Dellink et al. (2008) and apply respectively a normal distribution for α_i and β_i, and a restricted gamma distribution for s_i. The latter restriction assures that the regional shares sum to one. The values for the parameters characterising the distributions are also equivalent to Dellink et al. (2008). The most important parameter for studying the influence of fat tails on coalitional stability and performance is the global benefit parameter γ. We specify a base scenario in which a normal distribution is specified for γ, with its mean and standard deviation based on the meta-analysis presented in Tol (2009). Their values are respectively set to 120 $/tC and 148 $/tC.

---

[15] During the Monte Carlo sampling procedure 20,000 samples are generated from their respective probability density functions. The discount rate is not specified as a stochastic parameter, since raising the discount rate in the STACO model has an equivalent impact as a lower mean value of gamma (Weikard et al., 2006).

In this base scenario, climate change can have positive and negative economic impacts, i.e. the global benefit parameter can be respectively negative (positive climate change impacts) and positive (negative climate change impacts).[16] A negative (or zero) value for γ_m implies regions have no incentive to conduct any abatement. Optimal payoffs are therefore zero regardless of the coalition formed. Then all coalitions are trivial and stable by definition, as they do not commit to action. The reader should keep in mind that the SL increases when the percentage of negative samples in the Monte Carlo analysis increases. The level of the performance indicators will be affected when comparing the performance of coalitions across alternative distributional specifications using the SL concept. To handle this aspect of the numerical model, we use the term of strict stability likelihood (SSL), which is defined as the SL minus the fraction of samples with negative draws (where climate change impacts are positive) during the Monte Carlo procedure. Strict stability likelihood reflects the model specification where trivial coalitions are not counted as being stable. The performance indicators success and efficiency are adjusted accordingly by replacing the SL by the SSL in equations (5) and (6).

4. Fat tail scenarios

Debates regarding uncertainty about climate change impacts are generally focused on low-probability, high-impact scenarios located in the upper tail of the global benefit parameter. Weitzman (2009a, 2009b), for example, argues that society has an infinite expected loss from such scenarios, if the upper tail is sufficiently fat. He calls for a precautionary approach where policies are based on preventing the worst case scenario. As argued in the introduction, we take a different approach and explicitly assume the expected value of climate change impacts remains finite, thereby assuming that IAMs and Cost-Benefit analysis remain valid tools for policy analysis. Within these boundaries, we are interested in the consequences of varying degrees of uncertainty, represented by different scenarios that vary the distribution of the global benefit parameter γ, on coalition stability and performance. Discussions on 'fat-tailed' climate change uncertainty generally focus on the upper tail of the distribution of climate

---

[16] Note that regional benefit shares are bounded from below at zero. This precludes situations where some countries have positive damages, while others have negative damages. While this is clearly a limitation of the model that affects the outcomes of the simulations, Dellink et al. (2008) present a sensitivity analysis on the gamma distribution by replacing it with a normal distribution, which implies that regional shares can become negative. They find that the impact on the stability analysis is very limited. Moreover, given the focus of our paper on the high end of the damage function, it is not unreasonable to assume that damages will be positive in all regions.

change impacts (e.g. Tol, 2009). Therefore, we investigate the impacts of increasing the upper tail of the respective distribution on coalition formation.

Climate uncertainty (and risk) and the fatness of the upper tail are measured in this paper by means of the four distributional moments, i.e. mean, standard deviation, skewness and kurtosis. There is not a single moment that specifically determines the fatness of the tail; rather, a specific combination of the different moments increases the probability of very high damages. Increasing the probability of high climate impacts may stem from an increase in the mean of the global benefit parameter, spurring abatement efforts and potentially cooperative behaviour.[17] The standard deviation captures the degree of uncertainty, i.e. spread, in climate change impacts. An increase in the standard deviation makes both tails of the distribution are fatter, i.e. more extreme samples become more likely, including very high damages. The skewness measures how symmetric the uncertainty is. Fat tails are generally characterised by left- (or negatively) skewed distributions, implying that the majority of observations has a value larger than the mean. Decreasing the skewness implies that the probability of high-impact events increases, relative to low-impact events. Finally, the kurtosis measures how closely the distribution is centred around the mode. The higher the kurtosis, the more uniform predictions about climate change impacts are, i.e. there is a large peak in the distribution. Fat upper tails are therefore mainly captured by the mean, standard deviation and skewness of the distribution of the global benefit parameter, while the kurtosis measures the extent to which uniform (or dispersed) predictions have an impact on the stability and performance of coalitions. Various degrees of climate change risks are imposed in our simulation model by specifying a set of alternative scenarios varying in (i) the distributional form, and (ii) the level of the underlying distribution-describing parameters. Accordingly, the four moments of the distribution of the global benefit parameter systematically vary across the scenarios allowing for an identification of the impact of fat tails on coalition formation and performance.

Our base scenario assumes a normal distribution for global benefits. The normal distribution has the characteristic of being symmetric and unbounded. In the first set of scenarios we increase the expected value of climate benefits by increasing the mean of the normal distribution by the factors [1.5; 2; 3; 4; 5]. In a second set of scenarios we increase the standard deviation by the same factors, while keeping the mean fixed. The normal distribution can, however, not be used to test the influence of all four moments, as its skewness and

---

[17] However, increasing mean benefits also increases free-riding incentives.

kurtosis are constant. Therefore, we propose the use of three asymmetric distributions, respectively the gamma, beta, and double-sided exponential. For each of these distributions, we specify a reference scenario (as listed in Annex II) and then vary the parameters as described below.

Each of these distributions has specific convenient properties. The gamma distribution has a pre-defined lower bound of climate damages[18], but damages can still reach up to infinity. The gamma distribution is right-skewed and approaches a normal distribution for specific combinations of the mean and standard deviation. We propose another set of five scenarios where we use the same mean value as in the normal distribution, but increase the standard deviation by the same factors as in the normal case. This increases the fatness of the right tail, but also implies more draws closer to the lower bound. The latter is a direct consequence of fixing the mean at 120 $/tC. In other words, the standard deviation, skewness and kurtosis are increasing in these scenarios (see Annex II).

The beta distribution is very flexible and can accommodate both left- and right-skewed shapes. We define its shape through a lower and upper bound, a mean, and standard deviation parameter. Values for the lower bound, mean and standard deviation are again set at respectively -80, 120 and 148 $/tC. The upper bound is assumed to be at 675 $/tC in those scenarios (the 99th percentile in Tol, 2009). By increasing the mean, compared to the reference scenario, we gradually work towards a left-skewed distribution with a large probability of high climate change damages. In a separate set of scenarios we gradually increase the upper bound to 3375 $/tC, while keeping the other parameters constant. This results in a set of right-skewed distributions with increasing skewness and kurtosis.

Finally, we use a double-sided exponential function as applied in Dellink et al. (2008). By applying a double-sided distribution we can independently alter the shape of both tails of the distribution. We base the parameter values again on Tol (2009) and specify the 5% mass point at -45 $/tC and the 95% mass point at 410 $/tC. The mode of the distribution is set at 49 $/tC and there is a 35% probability of positive climate change impacts, i.e. negative benefits (γ) from mitigation. For the exponential distribution, we specify a set of alternative scenarios around its reference specification. First, we increase the upper bound and thus the upper tail.

---

[18] Following the meta-study by Tol (2009) the lower bound of climate damages is set at -80 $/tC in all scenarios, based on the 1st percentile for global damages.

Second, we reduce the probability of having observations below the mode and thereby increase the mass in the right tail.

In total this results in a set of 41 scenarios with varying distributional shapes. See Annex I for a full overview of input parameters and Annex II for a full overview of the resulting sample moments. In contrast to Dellink et al. (2008), this set of scenarios allows for a structured econometric analysis of increased climate change risk and the influence of the four distributional moments on the stability likelihood and other performance indicators.

5. Simulation results: the impact of fat-tailed distributions on the stability of coalitions

5.1. Simulation results in the no-learning and full-learning case

To become familiar with the results from the STACO model, we begin with a discussion of the outcomes for the no-learning case for the normal distribution, our base scenario. The results are presented in Table 1. Three specific coalitions are presented, the all singletons, the grand coalition, and the best performing coalition (BPC). The latter is defined as the stable coalition with the highest expected global payoff from abatement, i.e. the coalition that scores best on our indicator of success. Total annual abatement efforts and global payoff are lowest in the AS case and highest for the GC. Obviously, most abatement efforts take place in the regions with low marginal cost, i.e. the USA and China, and abatement efforts increase due to the cooperative behaviour of coalition members. The net present value of global payoff for the GC is highest, but the positive values in the incentives column denote that the coalition lacks internal stability. In fact, all regions have an incentive to change their announcement and leave the GC. Remember that the sum of individual incentives to free-ride equals the required transfers to keep a player aboard, while the associated gains from cooperation are the additional gain for the coalition from moving from 11 to 12 players. For very small coalitions, the additional gain from cooperation is relatively large (as joint abatement efforts are still modest) and existing members can compensate the potential new member for foregone free-rider profits through the transfer scheme. As the coalition becomes larger, free-rider incentives increase and simultaneously the additional gains from cooperation for existing coalition members decrease. At some point, the free-rider incentives are larger than the additional coalition gains and the coalition is no longer potentially internally stable.

The BPC is both internally and externally stable, since none of the regions has a positive incentive to change its membership. Amongst its five members - China, DAE, EET, India,

and the US - China conducts the most abatement. The transfer column shows that the US pays all other coalition members, but especially China, to conduct these abatement efforts and to keep the coalition internally stable. The US benefits from these additional coalition members through its high level of marginal benefits.[19]

Table 1: No learning outcomes base scenario (normal distribution)

Regions	All singletons (AS)		Grand coalition (GC)				Best performing coalition (BPC)				Membership
	Annual abatement % of BAU	NPV of payoff bln US$	Annual abatement % of BAU	NPV of payoff bln US$	Incentives bln US$	Transfers bln US$	Annual abatement % of BAU	NPV of payoff bln US$	Incentives bln US$	Transfers bln US$	
USA	0.016	4431	0.037	11957	3575	-3117	0.022	6214	-40	-2406	1
JPN	0.006	3820	0.023	10582	3164	-4071	0.006	7542	-769	0	0
EEC	0.013	4898	0.030	13164	3936	-5214	0.013	9842	-989	0	0
OOE	0.009	809	0.047	2378	711	220	0.009	1598	-79	0	0
EET	0.008	305	0.075	901	269	800	0.046	571	-4	225	1
FSU	0.011	1474	0.041	4214	1260	273	0.011	2932	-222	0	0
EEX	0.004	685	0.049	1982	593	952	0.004	1352	-61	0	0
CHN	0.028	1202	0.100	2800	837	7583	0.086	1641	-11	1614	1
IND	0.018	986	0.100	2800	837	1288	0.063	1728	-11	327	1
DAE	0.004	537	0.059	1569	469	835	0.032	989	-6	241	1
BRA	0.000	338	0.011	1026	307	-199	0.000	665	-13	0	0
ROW	0.010	1248	0.050	3568	1067	650	0.010	2480	-171	0	0
Total	0.014	20733	0.052	56941	17026	0	0.028	37554	-2376	0	5

In the full-learning case, stability depends on the resolution of uncertainty in each sample of the Monte Carlo procedure. Hence, the concept of SL plays an important role. Table 2 shows that the AS and the GC have a relatively low SL of 21 per cent. That is, for most realisations of the global benefit parameter these coalitions are unstable. The coalition with the highest stability likelihood (HSLC) has eight members and abates (on average over all draws) substantially more than the AS coalition. The number of coalition members in the BPC, the

---

[19] Whether such a set of international transfers is politically realistic remains to be seen, but there is no doubt that countries with high benefits have an economic incentive to finance mitigation in other countries.

most successful coalition, is larger than in the HSLC from a global perspective. That is, the increase in benefits from additional cooperation outweighs the marginal reduction in SL. Note that the BPC in Table 2 differs from the BPC in Table 1 due to the alternative treatment of uncertainty in determining membership and abatement efforts in both settings.

Table 2: Highest SL and best performing coalition in the full learning base scenario

	Abatement % of BAU	Payoff bln US$	Incentives bln US$	Transfers bln US$	Abatement % of BAU	Payoff bln US$	Incentives bln US$	Transfers bln US$
	AS: SL = 21.3%				GC: SL = 21.4%			
Global	0.012	31412	1818	0	0.044	90669	19599	0
	HSLC: SL = 27.1%				BPC: SL =.26.0%			
	(USA, EET, EEX, CHN, IND, DAE, BRA, ROW)				(all regions, except JPN, EEC)			
USA	0.022	12759	124	-4304	0.025	14644	1617	-5039
JPN	0.006	14370	-1558	0	0.006	17145	-2642	0
EEC	0.012	18896	-1872	0	0.012	22569	-3298	0
OOE	0.008	3158	-268	0	0.032	3103	436	101
EET	0.046	1072	51	493	0.051	1177	167	573
FSU	0.009	5648	-592	0	0.029	5252	784	-237
EEX	0.029	2325	136	541	0.033	2573	380	613
CHN	0.067	3402	162	2969	0.070	3858	511	3578
IND	0.059	3106	231	237	0.063	3491	518	320
DAE	0.036	1734	146	359	0.040	1940	313	449
BRA	0.007	1130	92	-111	0.008	1258	205	-196
ROW	0.031	3761	255	-184	0.034	4236	608	-162
Total	0.029	71360	-3092	0	0.034	81246	-400	0

The lower part of Table 2 presents the regional-specific abatement and payoff levels for the HSLC and BPC under our base scenario and full-learning. Reported values are mean values, i.e. averaged over the 20,000 draws in the Monte Carlo procedure. In the HSLC and the BPC, the USA again subsidises China and other developing countries to become a member of the coalition and undertake most abatement. Note that the incentives column in Table 2 clearly illustrates the concept of SL. Only in 27% (26%) of the realisations of the global benefit parameter the HSLC (BPC) is stable, in all other cases some regions can benefit from changing their announcement. On average, Japan and Europe have no incentive to become a member of the HSLC or the BPC due to their high marginal abatement costs, while the US

can on average benefit from stepping out of both coalitions.[20] All three regions have high marginal benefits from abatement, which makes them willing to fund abatement efforts in other countries. The relatively higher marginal abatement costs in Japan and Europe makes these regions less attractive coalition partners than the US. Including Japan or Europe into the coalition instead of the US would imply shifting additional abatement efforts to the other partners in the coalition. In the absence of transfers all three regions are somewhat unattractive coalition members, because their high marginal abatement efforts stimulate free-riding incentives while they would contribute little in terms of abatement efforts given their high marginal costs. It should be noted that coalitions where Japan or Europe are cooperating with a set of developing countries also have relatively high SL-values, but these perform worse in terms of global payoffs. The key message is that transfers allow the international agreement to be formed by almost any mixture of countries that comprises countries with high marginal benefits and countries with low abatement costs. Attractiveness of partners depends on the relation between marginal abatement benefits and costs.

5.2. The SL under alternative distributional forms

To get a feeling for how alternative distributional forms affect the SL, the left panel of Figure 1 depicts the samples for the global benefit parameter that have been used as an input in the reference case analysis for each distributional forms, i.e. normal, gamma, beta, and double exponential (see Annex I for specifications). The right panel presents the resulting distribution of the SL over all possible coalition structures. The level of the SL is comparable between the normal and the gamma distribution. The main difference between these two distributions is in the left tail, but as optimal abatement levels and payoffs are zero for any level of global benefits below zero, the exact shape of the left tail is not relevant for the coalition stability. The double exponential brings about significantly lower SL values than the normal case, whereas the beta distribution shows the highest SL levels. As mentioned in Section 3.2, the differences in the level of SL across all distributional forms are mainly caused by the share of negative draws for the global benefit parameter (i.e. positive climate change impacts) in our

---

[20] By definition, the incentive not to join a coalition (i.e. the incentive for a current free-rider) is the inverse of the incentive to free-ride for the enlarged coalition, where the same player is a member. The overarching idea of the 'incentive' column is that negative values indicate a contribution to stability, while positive values undermine stability.

Monte Carlo analysis. This inspires us to us the concept of Strict Stability Likelihood in Section 6. More interesting is the right tail of the SL distribution function: very few coalitions have a substantially higher SL than the minimum, i.e. the kurtosis of the distribution is very high. We investigate this result in detail in Section 6.2.

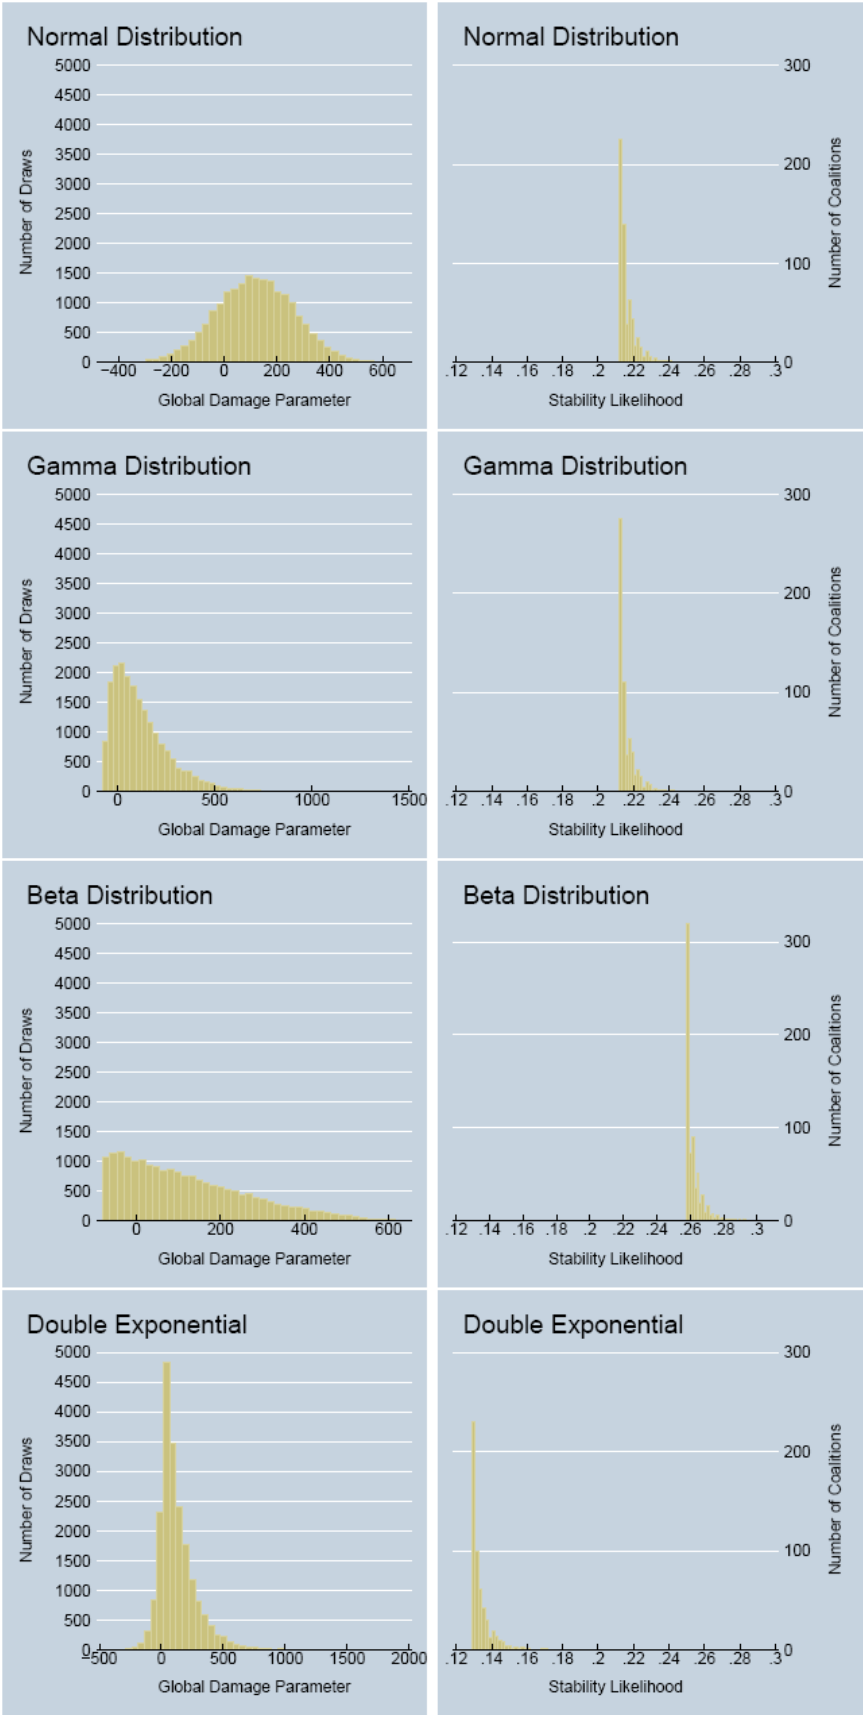## Figure 1: Inputs and SL for the reference scenarios

Table 3: Outcomes under full-learning for alternative distributional specifications

	Normal					**Gamma**				
	SL %	Payoff bln US$	Abatement MtC (%)	Incentives bln US$	Non-Members	SL %	Payoff bln US$	Abatement MtC (%)	Incentives bln US$	Non-Members
AS	21.3	31412	880(0.01)	1818		21.2	30264	800(0.01)	1715	
GC	21.4	90669	3260(0.04)	19599		21.3	85252	3020(0.04)	17858	
HSLC	27.1	71360	2140	-3092	JPN, EEC, OOE, FSU	27.1	69114	2020	-1968	USA, JPN, FSU
BPC	26	81246	2540	-400	JPN, EEC	25.8	76594	2330	-277	JPN, EEC
Negative draws (%)	21.3					21.2				
Expected outcome		64686	1950				61222	1780		
	Exponential					**Beta**				
	SL %	Payoff bln US$	Abatement MtC(%)	Incentives bln US$	Non-Members	SL %	Payoff bln US$	Abatement MtC(%)	Incentives bln US$	Non-Members
AS	12.9	31952	860(0.01)	1817		25.8	32008	820(0.01)	1826	
GC	13.1	89439	3290(0.04)	18542		25.9	90226	3000(0.04)	18875	
HSLC	19.5	72442	2190	-2090	USA, JPN, FSU	31.1	74924	2090	-2273	EEC, JPN, FSU
BPC	18.1	80345	2530	-395	JPN, EEC	30.3	81051	2330	-374	JPN, EEC
Negative draws (%)	12.9					25.8				
Expected outcome		64467	1940				64949	1790		

Table 3 presents SL, payoff and abatement levels for key coalitions as well as the share of negative draws for each of the four reference scenarios. In terms of coalition members, the BPC is continuously formed by the same countries across the alternative distributional forms, and is only lacking Japan and Europe as members. The HSLC is only the same for the gamma and double exponential scenario. For all coalitions, mean payoffs, reported in Table 3, are lowest for the gamma distribution. Since the mean and standard deviation for the normal and the gamma distribution are identical and the number of negative draws is very similar, the differences in mean payoffs can only be attributed to the upper tail, caused by variations in skewness and kurtosis of the gamma distribution. The right-skewness of the gamma distribution implies a smaller upper tail and suggests that fat tails may not complicate coalition formation (as the BPC is the same), but they have a positive effect on abatement

efforts and payoffs. The payoffs for the double exponential and beta distribution are comparable to the payoff for the normal base scenario, even though they have different skewness and kurtosis properties. This preliminary comparison of distributional forms is therefore insufficient to draw robust conclusions on the influence of changes in the upper tail on coalition formation.

The shaded rows in Table 3 reflect the expected outcomes for the specific scenario across all stable coalitions. Since it is not clear which stable coalitions will emerge in case of multiple stable coalitions, we define in each sample *m* the set of (undominated) stable coalitions, assign them an equal probability of occurrence and calculate mean payoff and abatement efforts for that sample. Expected outcomes reflect the average of these values over all 20,000 samples. Again the results for the double exponential and beta distribution are comparable to the payoff in the normal scenario, while for the gamma distribution, these values decrease by at least 5 per cent, which is in line with the results for individual coalitions discussed above.

6. Regression results: The impact of fat tails on coalition stability and performance

6.1. The regression setup

For each of the four distributional forms, multiple scenarios are included which differ in terms of a single input parameter relative to a reference scenario. By adjusting a single parameter, either the location or the shape of the distribution is adjusted, which directly affects one or more of the sample moments and thus the fatness of the tail.[21] The 41 scenarios are expected to have an impact on the SL, success, and efficiency of each of the 4084 possible coalition structures.[22] This provides us with a total number of 167,444 observations; one for each coalition in each of the 41 scenarios. We track the impact of these scenarios on the three performance indicators of interest (SL, success, and efficiency) using a regression analysis.

Two types of explanatory variables are included in the analysis. First, the scenario is summarised by the four sample moments (mean, standard deviation, skewness, and kurtosis)

---

[21] Note that our scenario selection mainly represents cases with high kurtosis when distributions are right-skewed. This affects the impact of the kurtosis on performance negatively.

[22] From the total of 4096 coalitions, all but one of the thirteen possible AS coalitions are excluded. We keep the AS coalition with no members. Using a definition for the AS with a single member would only affect the underlying incentives to change announcement, but not the performance of the coalition,

for the global benefit parameter as applied in the Monte Carlo analysis.[23] Second, we characterise the coalitions by a set of explanatory variables. It should be noted that the coalitions are constant across the scenarios and therefore do not reveal variations in their cost and benefit structures (parameters). Given that we only vary one dimension at a time, interactions between the coalitions and the scenarios are not directly taken into account. We chose to characterise coalitions in two alternative ways:

In Specification I, the coalition is described by specifying its member countries. The costs and benefits of abatement efforts for each country are described by three parameters, respectively a regional benefit share, s_i and two cost parameters, α_i and β_i. Given that these parameters are constant for each country across the scenarios, these operate like constants and can be summarised by country dummies. We introduce a set of dummies, which are set to 1 if a country is in the coalition and 0 otherwise. Every coalition structure can thus be created using country dummies, without inducing the incidental parameter problem (Lancaster, 2000).

These country dummies do not pick up interactions between particular countries and their impact on e.g. the stability of a coalition. We therefore present Specification II that focuses on the characteristics of the coalition rather than its member states. Explanatory variables are included describing the size of the coalition and the expected share of benefits covered by the members (i.e. the sum of the mean of s_i over the member states). Moreover, we stipulate that the stability, success, and efficiency of a coalition depend on the interaction between the member states. For example, countries with high benefits of abatement, but with high marginal costs, want to team up with low-cost countries to achieve their abatement efforts. Hence, we classify each coalition, except the AS, into a particular class.

We create 9 classes using the cost and benefit parameters of each country. Countries are grouped into three categories, respectively high, medium, and low, for both their costs and benefits (see Annex III). Then all possible interactions are taken into account, starting from those that are most beneficial for coalition formation (high benefit – low cost). A coalition is considered to be high benefit – low cost, if it contains at least one high-benefit and one low-cost country. In order for the categories to be mutually exclusive and prevent double-

---

[23] We prefer to work with the sample moments, since this allows obtaining the same moments for all distributions.

counting, the benefits are taken as point of departure. Consequently, not all nine categories will be represented in the analysis. Annex III provides a more detailed overview of the classification procedure. Given that the AS is not included in this structure, the associated parameters can be interpreted as performance relative to the AS.

The following subsections show our results for three different dependent variables, strict stability likelihood (SSL), strict success, and strict efficiency. For each of them we show two different sets of results: Specification I with country dummies and Specification II with coalition dummies. The regression setup is summarised by the following equation:

$$SSL_{sk} = \theta_1 * mean_s + \theta_2 * standard\ dev_s + \theta_3 * skewness_s + \theta_4 kurtosis_s + \theta_5 D_k + \varepsilon_{sk},$$

Whereas s = 1, …, 41, k = 1, …, 4084 and D contains either country or coalition dummies depending on the specification. All models are estimated with robust standard errors, clustered by coalition to correct for the panel specification.

6.2. Regression on Strict SL

In this section we track the sensitivity of the Strict SL (SSL) to varying specifications of the global benefit parameter, while controlling for the characteristics of the coalitions (Specification I and II). We do not use the SL, because it contains two counteracting effects: the level effect introduced by the negative draws and the effect we are interested in, arising from the collaboration of different regions (see Section 3.2). Therefore all regressions are conducted with the SSL to exclude influence from negative draws.[24]

As a reference case, we start off by looking at the performance of the AS coalition. It turns out its strict stability likelihood is zero across scenarios. In total, 192 coalitions (5%) never have a positive SSL, while 1781 coalitions (44%) have a positive SSL in each of the 41 scenarios. The maximum SSL observed is 7.2 percentage points, but 97% of the observations have a SSL smaller or equal to 2 percentage points. The limited spread in SSL is partially caused by the fact that the SSL is by definition reducing in the number of negative draws. Moreover, the principle of single deviations ensures that if a particular size is stable, all coalitions with one more or one less member will be unstable. Hence, only a limited number of coalitions will be stable at each draw during our simulation exercise.

---

[24] Negative draws imply that countries benefit from climate change and form coalitions without mitigating. The SSL is defined in the following way: SSL=SL - % of negative draws. Effectively, this implies that trivial coalitions are not counted as strictly stable.

The results for the SSL in Specification I, using the country dummies, are presented in Table 4. The SSL is increasing in the <u>mean</u> of the global benefit parameter. This implies that the additional benefits from cooperation outweigh the increase in free-rider incentives.[25] Increasing average benefits from abatement result in more stable coalitions. More specifically, the coefficient implies that increasing the mean of the global benefit parameter by one $/tC, will result in an increase of the SSL by 0.0002 percentage points. The coefficients for the other moments do not have a clear interpretation, since these are not related to the level of the global benefit parameter, but merely affect the shape of its distribution. Our result confirms the conclusion from Dellink et al. (2008) that higher expected benefits improve the chances for coalition formation.

A higher <u>standard deviation</u> amplifies both the probability of positive and negative climate change impacts, as it essentially flattens the distribution. The former implies a higher number of negative draws, where coalitions are stable by definition, but not strictly stable. This leads to a reduction of the SSL. A high standard deviation also increases the probability of higher benefits from abatement efforts, i.e. the tail becomes fatter, which may therefore increase the size of the stability and its stability (see Dellink et al. 2008). The net effect of the standard deviation turns out to be negative, while the impacts of the two opposing effects cannot be separated out.

A decrease in <u>skewness</u> results in more mass in the upper tail and highlights that a higher probability of large climate change impacts increases the SSL. We find that fat upper tails make coalitions more stable.[26] Thus, the answer to our original research question, whether fat tails stimulate coalition formation, is 'yes'. This effect comes on top of the effect that is more obvious, that higher mean benefits are good for coalition formation. Finally, an increase in the <u>kurtosis</u>, which puts more mass around the peak of the distribution, has a negative impact on the strict SL. When regions base their decisions on predictions of climate impacts that are conform and similar, the strict stability of coalitions reduces. Regions then seem to ignore contradicting estimates of climate change impacts, and outliers in the tails of the distribution appear to have a limited effect on the SSL.

---

[25] Cooperation between regions increases with higher marginal benefits from abatement as the additional benefits from mitigation can be used to finance transfers to keep low-marginal-cost countries within the coalition, thereby increasing the stability of the coalition. Simultaneously, the outside option for all countries increases with increasing marginal benefits, thus inducing additional free-riding incentives. The net effect is a priori ambiguous.

[26] A distribution with a fat upper tail is negatively skewed.

The results for the moments of the global benefit parameter are all significant at the 1% level and constant across our two specifications, highlighting the robustness of these results. We now discuss the effect the characteristics of the coalition have on the SSL. In both specifications the constant reflects the SSL of the AS coalition (as this is an 'empty' coalition which does not pick up the effects of the country or coalition dummies). In Specification I, the SSL decreases when countries with high benefits, especially Japan and the EEC, enter a coalition. This is caused by their high marginal costs and high marginal benefits, which make these two regions unattractive partners. High abatement costs are clearly reducing collaboration, as these regions will not contribute much to the joint abatement efforts. High benefits have a mixed effect: they allow for larger transfers to be distributed among coalition members, but they force all coalition members to undertake substantially more abatement. As marginal abatement costs increase more rapidly than marginal benefits (quadratic vs. linear), this tends to have a negative net effect on the incentives of other players to join the high-benefit countries. The USA (another high-benefit country) also has a negative impact on the strict SL by joining, although only significant at the 10% significance level. The low coefficient can be explained by its mixed attractiveness as a country with low marginal abatement costs, but high marginal benefits. In contrast, participation of the non-OECD regions increases the stability, primarily reflecting their low marginal costs, which make them attractive partners. As free-rider incentives increase with larger coalitions, this does not mean that coalitions with all 'attractive partners' are stable. Rather, coalitions consisting of one or two high-benefit countries together with one or two low abatement cost countries tend to be more successful in terms of strict stability likelihood.

This is further illustrated by Specification II in Table 4. SSL is increasing in the number of coalition members, while at the same time the SSL is decreasing in the share of benefits. The negative quadratic effects highlight that an additional member to the coalition does not tend to add as much to the SSL as the previous one. Moreover, as the share of benefits accumulates, the free-rider incentives increase. Finally, our set of explanatory variables, characterising the types of coalitions, show that high-benefit members want to join up with low-cost countries. This category has the highest positive coefficient. Members with medium or even low benefits are less attractive coalition members.

Comparing Specifications I and II, we see that in Specification I, the country dummies pick up two effects, which are separated in Specification II: (i) a higher share of benefits reaped by the coalition decreases stability (as free-rider incentives increase), but (ii) the presence of high-benefit countries in the coalition benefits stability (as more transfers are available).

Table 4: Regression results for strict stability likelihood

Specification I: distributional moments and country dummies			Specification II: distributional moments and coalition dummies		
	Coefficient	Std.error		Coefficient	Std.error
Mean	0.0002 ***	0.0000	**Mean**	0.0002 ***	0.0000
Standard dev.	-0.0002 ***	0.0000	**Standard dev.**	-0.0002 ***	0.0000
Skewness	-0.0088 ***	0.0003	**Skewness**	-0.0088 ***	0.0003
Kurtosis	-0.0017 ***	0.0000	**Kurtosis**	-0.0017 ***	0.0000
USA	-0.0261 *	0.0134	**Coalition size**	0.3810 ***	0.0228
Japan	-0.3952 ***	0.0133	**Coalition size sq.**	-0.0112 ***	0.0021
EEC	-0.1885 ***	0.0134	**Share of Benefits**	-2.6277 ***	0.0932
other OECD	-0.0957 ***	0.0134	**HB/HC**	0.1460 ***	0.0233
EET	0.2534 ***	0.0134	**HB/MC**	-0.0382	0.0488
FSU	-0.0336 **	0.0134	**HB/LC**	0.3396 ***	0.0532
EEX	0.1116 ***	0.0134	**MB/MC**	-0.5147 ***	0.0442
China	0.3156 ***	0.0134	**MC/LC**	-0.0939 **	0.0427
India	0.2503 ***	0.0134	**LB/MC**	-0.4493 ***	0.0496
DAE	0.2672 ***	0.0134			
Brazil	0.2457 ***	0.0134			
ROW	0.2089 ***	0.0134			
Constant	-0.0240	0.0266	**Constant**	-0.3424 ***	0.0208
R²	0.431			0.318	

Note: EEX: Energy exporting countries; EET: Eastern European countries; EEC: European Union; DAE: Dynamic Asian Economies; ROW: Rest of the World; * denotes 10% significance level; ** denotes 5%; *** 1%.

Note: ∑ Benefits: Share of benefits in the coalition; HB/HC: coalition has at least one high-benefit and one high-cost country; HB/MC: coalition has at least one high-benefit and one medium-cost country; HB/LC: coalition has at least one high-benefit and one low-cost country; MB/MC: coalition has at least one medium-benefit and one medium-cost country; LB/MC: coalition has at least one low-benefit and one medium-cost country; * denotes 10% significance level; ** denotes 5%; *** 1%.

6.3. Regression analyses on success and efficiency

The stability of a coalition, in the form of the SSL, provides limited information on the actual level of abatement conducted by that coalition or on its associated payoffs. Similarly, high payoffs from abatement, e.g. for the GC, turn out to be irrelevant if the coalition has a low stability likelihood. Therefore, we shed more light on the impact of uncertainty and fat tails

on the performance of a coalition by using our two indicators: Success and efficiency, as introduced in Section 2.2, and corrected using the SSL.

6.3.1. Performance indicator: Strict success

As discussed in Section 2, our success measure reflects the a priori expected stable payoff of a coalition. Again the AS coalition acts as our reference coalition. Given that it has a zero SSL, it also has zero strict success for every scenario. As revealed by Specification I in Table 5, increases in the <u>mean</u> of the global benefit parameter increase the success of a coalition through two channels. First, the result from Table 4 transfers that increases in mean benefits increase the stability of the coalition. Secondly, higher benefits from abatement increase expected global payoffs. In this case, the additional benefits from abatement outweigh the additional free-rider incentives that arise at the same time.

Increasing uncertainty regarding the impacts of climate change, as measured by the <u>standard deviation,</u> also increases the success of a coalition. This result is opposite from the impact of higher standard deviation on SSL. Like in the SSL case, a lower bound on success exists: when impacts of climate change become positive, countries will not undertake any abatement resulting in zero payoffs. This effect is ignored in the SSL measure. Increasing the mass in the upper tail spurs payoffs as benefits from abatement increase. The latter effect is further highlighted by the <u>skewness</u> parameter. Again, the negative parameter implies that when the upper tail becomes fatter the strict success of a coalition increases. Finally, parameter for the <u>kurtosis</u> reveals a similar effect as for the SSL regression. As the predictions about the impacts of climate change become more uniform (high spike in the density function), the strict success decreases. This is partially caused by the fact that our gamma distribution has a very strong peak with near zero impact from climate change. The results for skewness and kurtosis are comparable with the results from Table 4, indicating the dominance of the stability effect over the payoff effect from fat-tailed distributions.

Again, the results are stable across our two specifications. Both regressions show that the AS is unsuccessful. Moreover, coalitions including low-cost countries, like China, are likely to improve the success of the coalition, especially in combination with high-benefit countries. Additional members spur the success even more, but again the impact of an additional member is reducing with free-rider incentives increasing. The latter is reflected by the negative impact of the share of benefits on success.

Table 5: Regression results strict success

Specification I: distributional moments and country dummies			Specification II: distributional moments and coalition dummies		
	Coefficient	Std.error		Coefficient	Std.error
Mean	326.9 ***	7.9	Mean	326.9 ***	7.9
Standard dev.	95.3 ***	1.7	Standard dev.	95.3 ***	1.7
Skewness	-1408.1 ***	68.4	Skewness	-1408.1 ***	68.4
Kurtosis	-605.2 ***	9.7	Kurtosis	-605.2 ***	9.7
USA	-1485.6	2006.1	Coalition size	48049.8 ***	3646.5
Japan	-54792.0 ***	2001.0	Coalition size sq.	-825.9 **	355.4
EEC	-24310.0 ***	2005.5	Share of Benefits	-382430.1 ***	14672.0
other OECD	-10288.1 ***	2008.6	HB/HC	26248.2 ***	3901.4
EET	38586.5 ***	2012.2	HB/MC	2180.2	7807.3
FSU	-1394.2	2006.9	HB/LC	54059.4 ***	8506.8
EEX	19944.4 ***	2011.3	MB/MC	-66709.0 ***	6713.5
China	46133.6 ***	2012.5	MC/LC	-14048.6 **	6609.4
India	38353.2 ***	2012.2	LB/MC	-57892.5 ***	7096.1
DAE	40986.3 ***	2012.3			
Brazil	35431.8 ***	2012.2			
ROW	33435.7 ***	2011.9			
Constant	-107956.2 ***	5359.4	Constant	-131827.0 ***	3875.3
R²	0.372			0.315	
Note: see Table 4 for abbreviations; * denotes 10% significance level; ** denotes 5%; *** 1%.			Note: see Table 4 for abbreviations; * denotes 10% significance level; ** denotes 5%; *** 1%.		

6.3.2. Performance indicator: Strict efficiency

Most of the results for the strict success measure transfer to the regression for strict efficiency. The strict efficiency measure controls for the difference in levels of payoff across the various scenarios. It takes the payoff of respectively the GC and AS coalition as the upper and lower bound, measuring the extent to which the payoff of a coalition approaches the payoff of the GC coalition (i.e. it calculates the gains from cooperation as a percentage indicator). Like in the success measure, increases in payoffs and stability are both valued positively in this measure, but a correction is made for the difference in the size of the gap between AS and GC between different scenarios. Again, the AS coalition has a reference level of zero strict efficiency due to lack of strict stability and no gains from cooperation.

Table 6 reveals that again increases in the <u>mean</u> of the global benefit parameter has a positive effect on the regression results. As before, additional benefits from abatement stimulate abatement efforts in a coordinated fashion, making coalitions more efficient. The coefficient remains positive and significant, even when controlling for the number of countries in the

coalition in Specification II. Like in the SSL regression, the <u>standard deviation</u> has a negative impact on efficiency. The underlying intuition was already discussed in Section 6.2. The positive impact of higher standard deviation on success disappears with efficiency, as this impact accrues to all coalitions and therefore does not contribute to the relative gains from cooperation. The <u>skewness</u> again illustrates our main results. The larger the probability of high negative impacts of climate change, the more efficient coalitions will become. First, coalitions become more stable. Second, more countries are willing to join a coalition due to the expected increase in benefits from abatement. The additional free-rider incentives do not seem to be dominant in this process. Finally, the results for the <u>kurtosis</u> confirm the earlier insights. The lack of significance of the squared term for coalition size, which arises only in the regressions on efficiency, is explained by the fact that increases in the size of the coalition will always result in payoffs closer to the grand coalition.

Table 6: Regression result for strict efficiency

Specification I: distributional moments and country dummies			Specification II: distributional moments and coalition dummies		
	Coefficient	Std.error		Coefficient	Std.error
Mean	0.0001 ***	0.0000	**Mean**	0.0001 ***	0.0000
Standard dev.	-0.0001 ***	0.0000	**Standard dev.**	-0.0001 ***	0.0000
Skewness	-0.0033 ***	0.0002	**Skewness**	-0.0033 ***	0.0002
Kurtosis	-0.0009 ***	0.0000	**Kurtosis**	-0.0009 ***	0.0000
USA	-0.0006	0.0080	**Coalition size**	0.1319 ***	0.0157
Japan	-0.1855 ***	0.0079	**Coalition size sq.**	0.0013	0.0016
EEC other	-0.0736 ***	0.0080	**Share of Benefits**	-1.4044 ***	0.0610
OECD	-0.0283 ***	0.0080	**HB/HC**	0.1280 ***	0.0171
EET	0.1407 ***	0.0080	**HB/MC**	0.0486	0.0332
FSU	0.0026	0.0080	**HB/LC**	0.2483 ***	0.0363
EEX	0.0758 ***	0.0080	**MB/MC**	-0.1964 ***	0.0273
China	0.2176 ***	0.0080	**MC/LC**	-0.0384	0.0276
India	0.1470 ***	0.0080	**LB/MC**	-0.1633 ***	0.0257
DAE	0.1504 ***	0.0080			
Brazil	0.1266 ***	0.0080			
ROW	0.1270 ***	0.0080			
R²	0.396			0.291	
Note: see Table 4 for abbreviations; * denotes 10% significance level; ** denotes 5%; *** 1%.			Note: see Table 4 for abbreviations; * denotes 10% significance level; ** denotes 5%; *** 1%.		

6.4. Main insights from the regression analysis

Our regressions deliver a set of consistent results across the different moments of the stochastic benefit function (and specifications characterising the coalitions), as summarised in

Table 7. First, increasing expected benefits, i.e. a higher mean from abatement improves stability, size, and performance of coalitions. This confirms the insights from earlier studies (e.g. Dellink et al., 2008). Secondly, more uncertainty about the impacts of climate change (i.e. higher variance) reduces coordinative efforts, because more uncertainty also implies that positive impacts from climate change may arise. Higher variance may, however, improve the success of coalitions: The opportunity to avoid negative impacts of climate change by undertaking abatement efforts becomes more attractive, as the probability of high benefits from abatement rises. Thirdly, if the uncertainty is specifically directed to the upper tail, i.e. the fat tail of climate change (i.e. lower skewness), our analysis shows that countries are more willing to cooperate, which results in more stable and better performing coalitions. Fourthly, if predictions about climate change impacts become more uniform (i.e. higher kurtosis), our dependent variables show a decrease in stability. This result deserves attention since our scenarios mainly reveal a high kurtosis at the lower end of the distribution of the global benefit parameter. In policy terms, it suggests that the more we learn about climate change impacts (in terms of narrowing the bounds on possible impacts), the lower chances for stable high-performing coalitions. This mimics and to some extent qualifies the negative information effect found in more theoretical models such as Ulph (2004) and Kolstad and Ulph (2008). Fifthly, fat tails have a stronger impact of the stability of coalitions than on the performance of specific coalitions. Thus, they boost incentives to participate more than they boost ambition levels of countries.

Table 7: Overview of impact of the moments of the global benefit parameter

	Strict stability likelihood	Strict success	Strict efficiency
Mean	+	+	+
Standard deviation	−	+	−
Skewness	−	−	−
Kurtosis	−	−	−

The other explanatory variables in our regressions also provide some robust insights. First, characteristics of coalitions impact stability in multiple ways: (i) larger coalitions have better performance, but with each new member this effect is declining, (ii) the presence of high-benefit countries boosts participation as more transfers are available, and (iii) free-rider incentives increase with the ambition level of coalitions. This is in turn affected by the joint share of benefits reaped by the coalition, reducing the incentives to join a coalition.

7. Conclusions

In this paper we have examined the impact of increasing the fat tail of the distribution of climate change damages on the stability and performance of international mitigation agreements. More specifically, the impacts of variations in the fat tail on the formation of (stable) coalitions amongst twelve major world regions, and the resulting abatement efforts, were analysed within the applied game-theoretical STACO model. Variations in the fatness of the upper tail of the damage distribution were introduced by defining a range of alternative scenarios varying in the applied distributional form of climate change damages and by altering the parameters of these distributions. While the model is extremely stylised and not suited for in-depth analysis of mitigation policies in different countries, the main results can be put into the larger context of actual international negotiations on climate change.

First, even countries that do not participate in an international agreement implement policies to reduce greenhouse gas emissions. The USA is the typical example. On the one hand, it did not ratify the Kyoto Protocol, and efforts to implement a federal emission trading scheme (or carbon tax) have so far not been successful. On the other hand, it has adopted many policies at both state and federal level that effectively curb carbon emissions. OECD (2012) contains numerous examples of such mitigation actions by major polluting countries. These mitigation policies are also present in the STACO model where, even in the absence of an international climate agreement, most regions will undertake abatement efforts relative to their BAU scenario. Our main point here, which is confirmed by our simulation runs, is that the usual assumption in modelling baselines, that countries not in the coalition do nothing to reduce emissions, may well be false.

Within the STACO model, regions decide to abate based on the expected marginal costs and benefits of the required emission reduction activities. In the real world, it is impossible to verify what drives these unilateral abatement efforts. Emission reductions can be the result of co-benefits arising from, for example energy efficiency or energy security policies, but also genuine concerns about catastrophic damages may be a driving factor. Empirical verification of our first conjecture is further limited by the absence of a counter-factual against which to compare current emission trends. Nevertheless, the availability of numerous policy examples and an emerging literature on implicit carbon prices in different countries (e.g. Productivity

Commission, 2011) does point to some evidence that even in absence of a binding international agreement countries see scope for mitigation policies.

Secondly, by not being part of a larger coalition, regions are likely to conduct less abatement because they are unable to benefit from possible burden sharing schemes (i.e. monetary transfers) between members of the coalition. That is, while developing countries are in general more vulnerable to climate change than developed countries, the sheer size of economic assets at risk in developed countries imply that they have a large incentive to induce global mitigation action, and it is in their interest to assist developing countries financially. In this light, the US$100bn international climate fund agreed upon in Cancún in 2010, is rational. The STACO model treats such transfers, or flexibility mechanisms, in international climate agreements as a way to increase the efficiency of the international agreements by splitting the decision of who undertakes mitigation efforts from the decision who pays for it. Such flexibility mechanisms are already prominent in the Kyoto Protocol (incl. the Clean Development Mechanism and Joint Implementation) and are likely to be scaled up in future international regimes.

Given that the level of mitigation is less than socially optimal in the unilateral case, our analysis shows that an ambitious international agreement can still bring global benefits by further reducing climate risks. The fat-tail scenarios show that countries are willing to form coalitions to reduce the risk of irreversible climate damages, i.e. shave off the fat tail. We find that both, higher expected benefits and fatter upper tails of climate impacts, increase the strict stability likelihood of an international climate agreement. Hence, the probability of a stable agreement increases in the possibility of low-probability, high-impact climate change scenarios. Simultaneously, the expected payoffs are increasing in the likelihood of very high climate impacts. This spurs abatement efforts both within and outside coalitions. Thus, fatter tails lead to 'fatter' climate coalitions in two ways: they increase the performance of existing coalitions and they make larger coalitions stable. These observations are in line with many countries expressing the need to start (joint) abatement efforts soon.

Thirdly, there remain significant difficulties in striking an ambitious international agreement with broad participation. Since Barrett (2003), it is generally clear that international agreements may be broad or deep, but not both. In other words, to get wide participation in an agreement the stakes (and hence free-rider incentives) should be low, while ambitious

agreements will only attract a few members. This is clearly illustrated by the estimated parabolic relationship between coalition size and the stability and success of a coalition. Our payoff functions show that the stakes are quite high (i.e. mitigation costs are substantial), and hence the free-riding incentives make it hard to form a stable coalition. Unfortunately, the current state of negotiations seems to confirm this point forcefully. While the number of countries that have pledged emission reductions in the Copenhagen Accord, and confirmed them in the Cancún Agreements and at the Conference of Parties in Durban, is substantial, the joint pledges are generally seen as insufficient to be on a least-cost pathway to keep global average temperature increases limited to 2°C (UNEP, 2010; and confirmed by subsequent updates in 2011 and 2012).

There is no direct relationship between the 2°C scenario and our analysis of fat-tails, since the former is more a political goal than the result of a cost-benefit type of analysis used in the STACO model and other integrated assessment models. However, the principles of the Nash equilibrium underlying our simulation runs support the emergence of a pledge-and-review type international architecture, where pledges are based on voluntary national announcements rather than legally binding reduction targets. Our analysis shows that increasing the fatness of the upper tail of climate damages makes it more likely that the benefits from abatement will match with the marginal costs, thereby providing more scope for a larger coalition.

Although there are clear links with actual policy developments, it goes without saying that the numerical results depend on the specification of the model. Dellink (2011) deals with this extensively using a deterministic approach. The stochastic approach adopted in this paper is a major improvement, given the pessimistic conclusion of Dellink (2011) that "the most uncertain assumption [on regional benefits] is also one of the most critical". Our stochastic framework makes the large uncertainties on benefits from abatement efforts explicit and traces its impacts on various indicators of the performance of coalitions.

References

Anthoff, D., Tol, R., 2010. Climate policy under fat-tailed risk: An application of FUND. ESRI Working Paper 348, Dublin.

Barrett, S., 1994. Self-enforcing international environmental agreements. Oxford Economic Papers 46, 878 – 894.

Barrett, S., 2003. Environment and statecraft: The strategy of environmental treaty making. Oxford University Press, Oxford.

Carraro, C., Siniscalco, D., 1993. Strategies for the international protection of the environment. Journal of Public Economics 52(3), 309 – 328.

Carraro, C., Eyckmans, J., Finus, M., 2006. Optimal transfers and participation decisions in international environmental agreements. The Review of International Organizations 1(4), 379-96.

Chander, P. and Tulkens, H. 1995. A core-theoretic solution for the design of cooperative agreements on transfrontier pollution. International tax and public finance 2(2) 279-293.

d'Aspremont, C., Jaquemin, A., Gabszewicz, J.J., Weymark, J.A., 1983. On the stability of collusive price leadership. Canadian Journal of Economics 16(1), 17-25.

Dellink, R.B., 2011. Drivers of stability of climate coalitions in the STACO model. Climate Change Economics 2(2), 105-128.

Dellink, R.B., Nagashima, M., van Ierland, E.C., Hendrix, E.M.T., Sáiz, E. and Weikard, H.-P., 2009. STACO Technical Document 2: Model description and calibratrion of STACO 2.1, mimeo, Wageningen University, available at www.enr.wur.nl/uk/staco.

Dellink, R.B., Finus, M., Olieman, N., 2008. The stability likelihood of an international climate agreement. Environmental and Resource Economics 39, 357–377.

Dellink, R.B. and Finus, M. ,2012. Uncertainty and climate treaties: Does ignorance pay. forthcoming in Resource and Energy Economics.

Dietz, S., 2009. High impact, low probability? An empirical analysis of risk in the economics of climate change. Grantham Research Institute on Climate Change Working Paper 9, London.

Ellerman, A.D., Decaux, A., 1998. Analysis of post-Kyoto CO2 emissions trading using marginal abatement curves, Joint program on the science and policy of global change. Report No. 40, MIT, Cambridge, MA.

Finus, M., van Ierland, E.C., Dellink, R.B., 2006. Stability of climate coalitions in a cartel formation game. Economics of Governance 7(3), 271-91.

Ikefuji, M., R.J.A. Laeven, J.R. Magnus and Muris, C. 2010. Expected utility and catastrophic risk in a stochastic economy-climate model. CentER working paper No. 2010-122, University of Tilburg, Netherlands.

IPCC, 2007. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, M.L. Parry, O.F. Canziani, J.P. Palutikof, P.J. van der Linden and C.E. Hanson (eds). Cambridge University Press, Cambridge, United Kingdom.

Kolstad, C.D., 2007. Systematic uncertainty in self-enforcing international environmental agreements. Journal of Environmental Economics and Managament 53(1), 68-79.

Kolstad, C.D., Ulph, A., 2008. Learning and international environmental agreements. Climatic Change 89(1-2), 125-41.

Kolstad, C.D., Ulph, A., 2011. Uncertainty, Learning and Heterogeneity in International Environmental Agreements. Environmental and Resource Economics 50(3), 389-403.

Lancaster, T. 2000. The incidental parameter problem since 1948. Journal of Econometrics, 95(2) 391-413.

Na S., Shin, H.S., 1998. International environmental agreements under uncertainty. Oxford Economic Papers 50,173–185.

Nagashima, M., Dellink, R., van Ierland, E., Weikard, H.P., 2009. Stability of international climate coalitions – A comparison of transfer schemes, Ecological Economics 68(5), 1476-87.

Nordhaus, W.D., Yang, Z., 1996. A regional dynamic general-equilibrium model of alternative climate-change strategies. American Economic Review 86(4), 741-765.

Nordhaus, W., 2009. An analysis of the dismal theorem.
Cowles Foundation Discussion Paper, Department of Economics, Yale University, New Haven.

OECD, 2012. Environmental Outlook to 2050: The consequences of inaction. Organisation for Economic Co-operation and Development, Paris.

Olieman, N.J., Hendrix, E.M.T., 2006. Stability likelihood of coalitions in a two-stage cartel game: an estimation method. European Journal of Operational Research 174(1), 333-48.

Paltsev, S., Reilly, J.M., Jacoby, H.D., Eckaus, R.S., McFarland, J., Sarofim, M., Asadoorian, M., Babiker, M., 2005. The MIT emissions prediction and policy analysis (EPPA) model: Version 4. MIT Joint Program Report Series 125, Cambridge, MA.

Productivity Commission 2011, Carbon Emission Policies in Key Economies, Research Report, Canberra.

Tol, R.S.J., 2009. The economic effects of climate change. The Journal of Economic Perspectives 23, 29–51(23).

Ulph, A., 2004. Stable international environmental agreements with a stock pollutant, uncertainty and learning. The Journal of Risk and Uncertainty 29(1), 53-73.

UNEP, 2010. The emissions gap report: Are the Copenhagen Accord pledges sufficient to limit global warming to 2°C or 1.5°C?. UNEP, Geneva.

Weikard H.-P., Finus M., Altamirano-Cabrera J. C., 2006. The impact of surplus sharing on the stability of international climate coalitions. Oxford Economic Papers 58, 209–232.

Weikard, H.-P., 2009. Cartel stability under optimal sharing rule, The Manchester School 77(5), 575-93.

Weitzman, M., 2009a. Additive damages, fat-tailed climate dynamics and uncertain discounting. Harvard Economics Discussion Paper 26, Cambridge, MA.

Weitzman, M., 2009b. On modelling and interpreting the economics of catastrophic climate change. Review of Economics and Statistics 91(1), 1-19.

Annex I: Distributional forms for the global benefit parameter

Scenarios		Factor	Mean ($/tC)	Std. dev. ($/tC)
Normal	Base scenario		120	148
	Mean	1.5	180	148
		2	240	148
		3	360	148
		4	480	148
		5	600	148
	Sigma	1.5	120	222
		2	120	296
		3	120	444
		4	120	592
		5	120	740

		Factor	Lower bound ($/tC)	Mean ($/tC)	Std dev. ($/tC)
Gamma	Reference		-80	120	148
	Sigma	1.5	-80	120	222
		2	-80	120	296
		3	-80	120	444
		4	-80	120	592
		5	-80	120	740

		Factor	Lower bound ($/tC)	Mean ($/tC)	Std dev. ($/tC)	Upper bound ($/tC)
Beta	Reference		-80	120	148	675
	Mean	1.5	-80	180	148	675
		2	-80	240	148	675
		3	-80	360	148	675
		4	-80	480	148	675
	Upper bound	1.5	-80	120	148	1012.5
		2	-80	120	148	1350
		3	-80	120	148	2025
		4	-80	120	148	2700
		5	-80	120	148	3375

		Factor	5th % ($/tC)	Mode ($/tC)	PR<Mode	95th % ($/tC)
Exponential	Reference		-45	49	0,3456	410
	Pr < Mode	0.2	-45	49	0,31104	410
		0.3	-45	49	0,27648	410
		0.4	-45	49	0,24192	410
		0.5	-45	49	0,20736	410
		0.6	-45	49	0,1728	410
		0.7	-45	49	0,13824	410
		0.8	-45	49	0,10368	410
		0.9	-45	49	0,06912	410
	95% point	1.5	-45	49	0,3456	615
		2	-45	49	0,3456	820
		3	-45	49	0,3456	1230
		4	-45	49	0,3456	1640
		5	-45	49	0,3456	2050

Annex II: Sample moment for the global benefit parameter

Scenarios		Factor	Mean ($/tC)	St. dev ($/tC)	Skewn.	Kurt.	% < 0	5th % ($/tC)	95th % ($/tC)
Normal	Base scenario		118	148	-0.01	-0.02	21	-124	360
	Mean	1.5	179	147	0.02	0.00	11	-64	423
		2	239	148	-0.02	-0.07	5	-4	479
		3	359	148	0.04	-0.02	0	120	603
		4	480	148	-0.02	0.05	0	235	723
		5	600	148	0.00	-0.05	0	359	843
	Sigma	1.5	118	223	-0.00	-0.00	30	-248	487
		2	117	295	-0.00	-0.03	34	-368	603
		3	120	446	0.03	-0.07	40	-609	861
		4	134	595	0.01	0.03	42	-854	1104
		5	122	742	0.01	-0.00	44	-1104	1347
Gamma	Reference		118	146	1.45	3.09	21	-48	401
	Sigma	1.5	119	221	2.21	7.11	37	-74	564
		2	120	298	3.00	13.13	50	-79	714
		3	125	453	4.39	28.08	64	-80	986
		4	122	600	5.95	51.70	74	-80	1072
		5	127	785	8.12	105.80	80	-80	1107
Beta	Reference		119	148	0.78	-0.08	26	-64	408
	Mean	1.5	180	148	0.46	-0.48	11	-31	449
		2	239	148	0.20	-0.66	4	8	493
		3	361	148	-0.23	-0.64	1	105	588
		4	479	148	-0.81	-0.03	0	190	660
	Upper bound	1.5	119	148	1.03	0.84	24	-59	412
		2	120	149	1.18	1.51	23	-56	412
		3	120	148	1.21	1.69	22	-54	411
		4	120	147	1.28	2.13	22	-52	411
		5	120	149	1.38	2.51	22	-51	416
Exponential	Reference		125	149	1.65	5.42	13	-48	408
	Pr < Mode	0.2	143	175	-1.36	13.74	6	-58	410
		0.3	147	148	0.46	4.87	7	-48	416
		0.4	145	143	1.02	4.91	8	-45	406
		0.5	142	144	1.23	4.98	9	-45	407
		0.6	140	145	1.44	5.61	10	-44	411
		0.7	136	145	1.37	3.90	10	-43	416
		0.8	133	145	1.41	4.28	11	-45	410
		0.9	127	147	1.60	5.08	12	-44	405
	95% point	1.5	174	220	1.93	5.63	13	-45	612
		2	229	295	2.06	6.30	12	-44	817
		3	332	440	2.13	6.48	13	-44	1211
		4	437	595	2.20	6.88	13	-45	1645
		5	540	742	2.24	7.13	13	-45	2064

Annex III

The countries are classified with respect to benefit and cost characteristics:

	High benefit	Medium benefit	Low benefit
High cost	JPN		BRA
Medium cost	EEC	IND, ROW	OOE, EET, EEC, DAE
Low cost	USA	FSU, CHN	

Then we sort the 4083 coalitions (all except the AS which is used as the reference coalition in the regressions) in nine different classes, based on contribution to effective coalitions, i.e. starting with high benefits and low costs and working through the columns of the matrix:

	High benefit (HB)	Medium benefit (MB)	Low benefit (LB)
Low cost (LC)	3199 coalitions	382	0
Medium cost (MC)	381	94	26
High cost (HC)	1	0	0

HB/LC specifies a coalition with at least one high-benefit and one low-cost country. HB/MC describes a coalition with at least one high-benefit and one medium-cost country and no LC country in the coalition. HB/HC describes a coalition that includes at least one high-benefit and one high-cost country, and no MC or LC countries. Using the same procedure, we classified medium-benefit and low-benefit coalitions. Therefore, the nine categories are mutually exclusive and coalitions not double-counted.

If a coalition consists of for instance JPN (HB/HC) and OOE (LB/MC) and CHN (MC/LC), it would be classified as HB/LC, because JPN has high benefits and CHN low costs and it thus has the most favourable combination for effective cooperation. It does not matter how many countries are in a coalition, for the classification we just take into account the most "extreme" members.