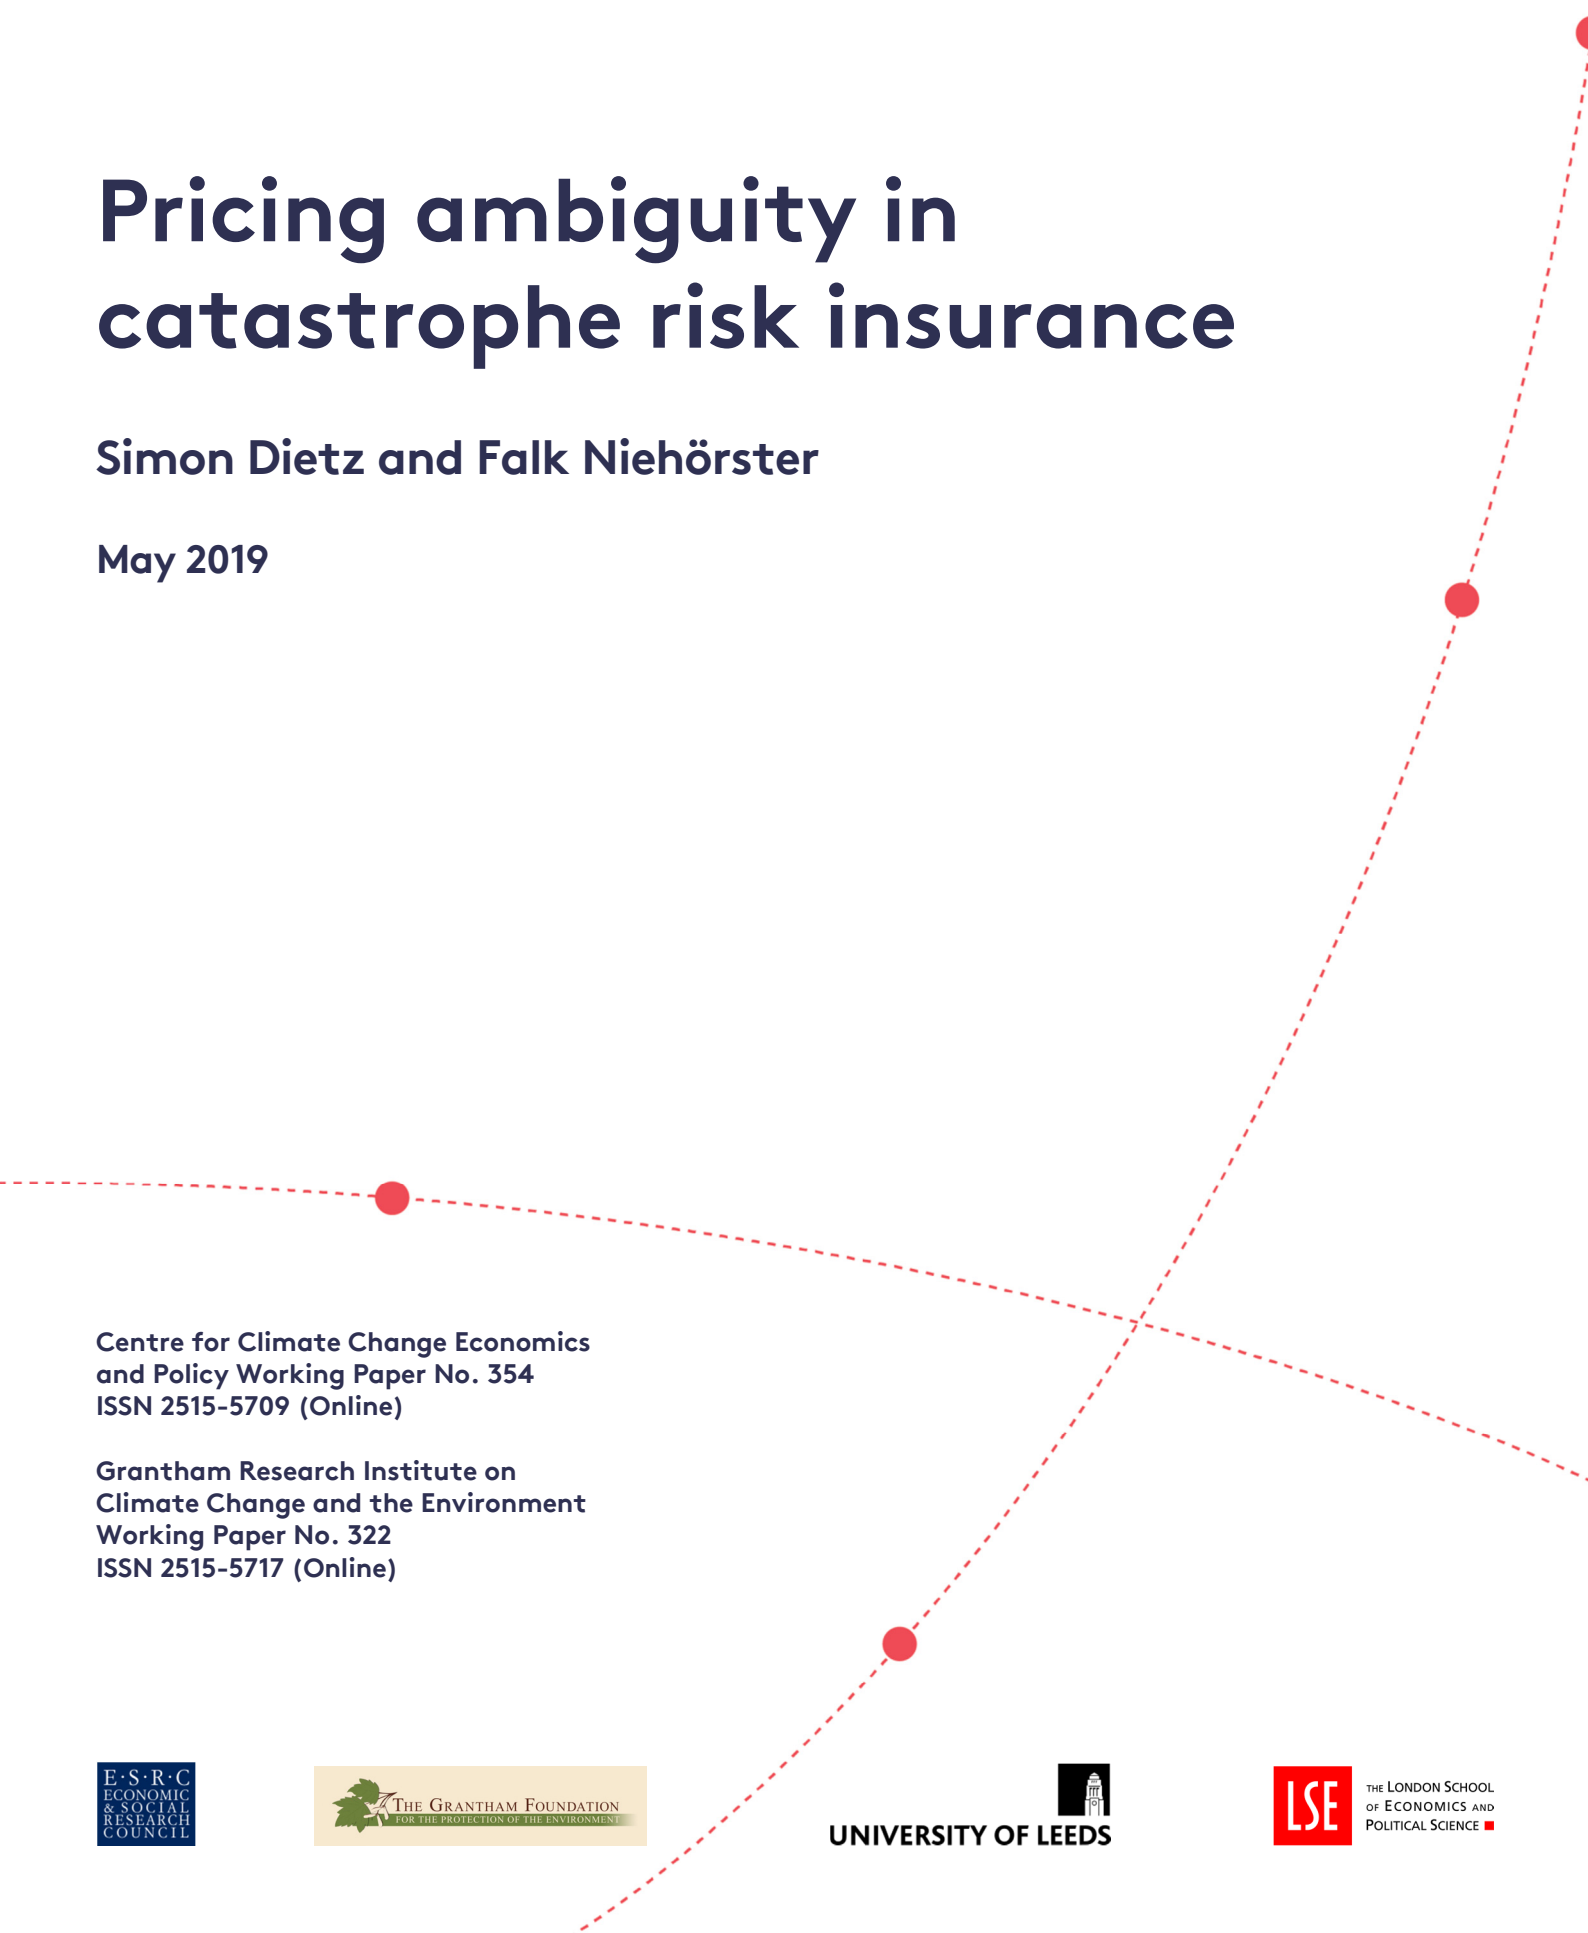


Pricing ambiguity in catastrophe risk insurance

Simon Dietz and Falk Nihörster

May 2019



Centre for Climate Change Economics
and Policy Working Paper No. 354
ISSN 2515-5709 (Online)

Grantham Research Institute on
Climate Change and the Environment
Working Paper No. 322
ISSN 2515-5717 (Online)

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council. Its third phase started in October 2018 with seven projects:

1. Low-carbon, climate-resilient cities
2. Sustainable infrastructure finance
3. Low-carbon industrial strategies in challenging contexts
4. Integrating climate and development policies for 'climate compatible development'
5. Competitiveness in the low-carbon economy
6. Incentives for behaviour change
7. Climate information for adaptation

More information about CCCEP is available at www.cccep.ac.uk

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training. The Institute is funded by the Grantham Foundation for the Protection of the Environment and a number of other sources. It has six research themes:

1. Sustainable development
2. Finance, investment and insurance
3. Changing behaviours
4. Growth and innovation
5. Policy design and evaluation
6. Governance and legislation

More information about the Grantham Research Institute is available at www.lse.ac.uk/GranthamInstitute

Suggested citation:

Dietz S and Niehörster F (2019) *Pricing ambiguity in catastrophe risk insurance*. Centre for Climate Change Economics and Policy Working Paper 354/Grantham Research Institute on Climate Change and the Environment Working Paper 322. London: London School of Economics and Political Science

Pricing ambiguity in catastrophe risk insurance

Simon Dietz^{*†} and Falk Nießhörster[‡]

May 23, 2019

Abstract

Ambiguity about the probability of loss is a salient feature of catastrophe risk insurance. Evidence shows that insurers charge higher premiums under ambiguity, but that they rely on simple heuristics to do so, rather than being able to turn to pricing tools that formally link ambiguity with the insurer's underlying economic objective. In this paper, we apply an α -maxmin model of insurance pricing to two catastrophe model data sets relating to hurricane risk. The pricing model considers an insurer who maximises expected profit, but is sensitive to how ambiguity affects its risk of ruin. We estimate ambiguity loads and show how these depend on the insurer's attitude to ambiguity, α . We also compare these results with those derived from applying model blending techniques that have recently gained popularity in the actuarial profession, and show that model blending can lead to the counter-intuitive result that the insurer prices catastrophe risk contracts as if it *seeks* ambiguity.

Keywords: ambiguity, catastrophe modelling, insurance, model blending, natural disasters

^{*}Grantham Research Institute on Climate Change and the Environment, and Department of Geography and Environment, London School of Economics and Political Science, London, UK

[†]Oxford Martin School and Smith School of Enterprise and the Environment, University of Oxford, Oxford, UK

[‡]DEVK RE, Cologne, Germany

We are very grateful for the input of Richard Bradley and seminar participants at ETH Zürich. We would like to acknowledge the financial support of the ESRC Centre for Climate Change Economics and Policy and the Grantham Foundation for the Protection of the Environment. The usual disclaimer applies.

Email for correspondence: s.dietz@lse.ac.uk.

1 Introduction

Ambiguity, defined as uncertainty about the relative likelihood of events (Ellsberg, 1961), is a common feature of insurance (Hogarth and Kunreuther, 1989). Ambiguity is particularly likely to affect catastrophe risk insurance, where historical loss data are limited. Hurricane risk provides a good motivating example of the issues. A number of insurers failed in the aftermath of Hurricane Andrew, which made landfall as a category four hurricane in south Florida in August 1992, principally because they had been relying on loss data from the previous two decades for pricing, reserving and so on (Calder et al., 2012). But 1971-1991 (or thereabouts) was an historically quiet period for Atlantic hurricane activity (Landsea et al., 1996),¹ as well as being a period during which insured values and other aspects of the state's vulnerability and insurers' exposure were changing. Partly as a consequence of Andrew, probabilistic catastrophe models have become widely used in the insurance industry as a means of estimating catastrophe risks. However, while catastrophe models may provide a better basis for estimating catastrophe risk than relying on the historical record alone, such models tend not to provide unambiguous loss probabilities. For example, Mao (2014) recently reported widely differing estimates of insured losses from a 1/250-year Florida hurricane, ranging from \$54 billion to \$151bn, with the three leading commercial catastrophe model providers disagreeing on both the mean loss and the range.²

Ambiguity tends to result in higher premiums than those charged for equivalent, unambiguous risks. This was a feature of industry practice that was widely suspected, but difficult to isolate from other factors tending to raise premiums. Therefore a series of studies used experimental and survey methods to establish, in controlled conditions, the practice of loading pre-

¹The average number of intense hurricanes (categories 3-5) forming in the Atlantic basin between 1944 and 1991 was 2.2 per year, whereas between 1971 and 1991 it was 1.5, roughly 0.5 standard deviations below the 1944-1991 mean.

²In this paper, we will not draw a distinction between uncertainty about the probability of a known loss and uncertainty about the loss conditional on an event of known probability (cf. Kunreuther et al., 1993). Both are regarded here as instances of ambiguity and can be rendered equivalent in terms of non-unique estimates of the loss distribution.

miums under ambiguity (Hogarth and Kunreuther, 1989, 1992; Kunreuther et al., 1993, 1995; Cabantous, 2007; Kunreuther and Michel-Kerjan, 2009; Cabantous et al., 2011). The participants in these studies, mostly insurance professionals (including actuaries and underwriters), were asked to price hypothetical contracts, with the existence of ambiguity about losses being a treatment. The results consistently showed that higher prices were quoted under ambiguity compared with when a precise probability estimate was available.³ This insurer ambiguity aversion is consistent with a broader body of evidence on individual decision-making under ambiguity (reviewed by Machina and Siniscalchi, 2014), initiated by Ellsberg’s thought experiments about betting on balls being drawn from urns (Ellsberg, 1961).

While it is well known that ambiguity contributes to higher insurance premiums and accompanying problems of availability and coverage, the economic rationale for this has received less attention. Hogarth and Kunreuther (1989) showed that a positive ambiguity load is consistent with a risk-neutral insurer who maximises expected profit, if the insurer’s beliefs about the probability of loss are formed according to the psychological model of Einhorn and Hogarth (1985). In this model, the insurer’s probability estimate is distorted such that, for low-probability events like catastrophe risks, greater weight is placed on higher probabilities of loss. That is, the insurer forms pessimistic beliefs when insuring catastrophe risks. A basic result of decision theory is that expected utility maximisation is inconsistent with ambiguity aversion (Ellsberg, 1961). However, Berger and Kunreuther (1994) showed that an ambiguity load could be consistent with a risk-averse insurer who maximises expected utility, if (and only if) the effect of ambiguity is to increase the correlation between risks in the insurer’s portfolio.⁴ They also showed, again under special assumptions about the nature of

³Cabantous (2007) provided some evidence that ambiguity increased premiums more when it was defined as the existence of conflicting, precise estimates, rather than a consensual, imprecise estimate. Our case studies in this paper are easiest to interpret as comprising conflicting, precise estimates.

⁴In particular, they showed that the premium per risk was an increasing function of the number of identical risks insured, when the loss probability was uncertain and either uniform or discrete distributed with an expected value of one half. This illustrated that under ambiguity the risks became correlated.

the ambiguity,⁵ that an ambiguity load is consistent with a ‘safety-first’ model of an expected-profit-maximising insurer who must satisfy an insolvency/ruin/survival constraint in the tradition of Stone (1973). Moreover they showed that this safety-first model better explains other results in the experimental/survey literature than does the expected utility model. What makes this approach attractive is that survival constraints are widely used in the insurance industry and have become enshrined in some regulations such as the European Union’s Solvency II Directive.

Dietz and Walker (2019) have recently proposed an alternative α -maxmin representation of the insurer’s decision problem under ambiguity, which is also based on an expected profit maximiser facing a survival constraint. The insurer’s attitude to ambiguity potentially affects the premium through the amount of capital it must hold against the risk of ruin. An ambiguity-averse insurer places more weight on higher probabilities of ruin, therefore holds more capital, and charges a higher capital load on the premium if the new contract increases the probability of ruin (*vice versa* for an ambiguity-seeking insurer). This is also a form of pessimism. The advantages of the model in Dietz and Walker (2019) are that it is consistent with ambiguity loading under quite general conditions whereby a new contract increases the ambiguity of the insurer’s portfolio, and that, drawing on recent advances in decision theory (especially Ghirardato et al., 2004), it captures the insurer’s attitude to ambiguity via an easily interpretable parameter, α , that is clearly separated from the insurer’s beliefs about the probabilities of loss.

The actuarial literature that directly engages with ambiguity would also appear to be sparse. Two areas are particularly relevant to the current setting. The first is the development of premium principles. In this regard, Pichler (2014) observes that all well-known premium principles assume a unique, stable loss distribution. One or two of these principles take ambiguity into account indirectly through the practice of distorting the loss distribution, placing more weight on higher loss probabilities and therefore

⁵Specifically their results were obtained using a discrete parameter distribution, assigning a probability of 0.9 on a loss probability of 0, and 0.1 on a loss probability of 1.

introducing pessimistic beliefs like the Einhorn and Hogarth (1985) model. The second is the development of techniques, primarily by practitioners, to mix multiple, conflicting probability estimates into a single estimate that can be inputted into a standard premium principle such as the expected value principle. These have come to be referred to as catastrophe model ‘blending’ (Calder et al., 2012). The interesting feature of these techniques, from our point of view, is that their development has been based on satisfying certain mathematical properties and generally without attention to insurers’ fundamental economic objective.

In view of the relatively limited body of theory on how ambiguity affects insurance supply,⁶ it is perhaps unsurprising that insurance actuaries and underwriters appear to rely on simple rules of thumb in practice when loading premiums under ambiguity. Hogarth and Kunreuther (1992) provided some structured evidence of this from their mail survey of actuaries. A subset of respondents provided written comments giving insight into their decision processes. Most of these responses indicated that actuaries first anchored the premium on the expected loss, and the majority of these would then, when informed the probability of loss was ambiguous, explicitly or implicitly apply an *ad hoc* adjustment factor or multiplier (e.g. increase the premium by 25%). Anecdotally such practice is widespread.

The purpose of this paper is twofold. The first is to apply the model in Dietz and Walker (2019) to two data sets, where conflicting catastrophe models create ambiguity about the probability of insured losses. The two data sets both relate to hurricane risk in the Atlantic basin. In one, Florida property, ambiguity stems from a set of (simple) hurricane models yielding different estimates of hurricane activity. In the other, Dominica property, we assume a unique estimate of the probability of hurricane activity, but there is nonetheless ambiguity, which stems from different vulnerability functions describing a property portfolio of uncertain quality/resilience. These applications enable us to demonstrate the practical use of an economic premium

⁶This can be compared with the relatively more extensive literature on ambiguity aversion on the demand side, e.g. (Snow, 2011; Alary et al., 2013; Gollier, 2014; Berger, 2016).

principle that treats ambiguity explicitly and thereby quantifies the ambiguity load as a function of the insurer’s attitude to ambiguity. We argue the application of this principle has the capacity to improve on *ad hoc* adjustments, even if it is unable to avoid the need for them altogether. The second purpose of the paper is to compare the results of this with the premiums that would be quoted by an insurer, who is also an expected profit maximiser operating under a survival constraint, but where the unique probability of loss is derived from the application of popular model blending techniques. The comparison yields the important and counter-intuitive result that these model blending techniques can be inconsistent with insurer ambiguity aversion, instead implying the insurer is ambiguity-seeking. This will not always be the case and it is beyond the scope of our paper to demonstrate how often it will be the case. We merely set out to demonstrate that it can happen, by providing two plausible examples.

The remainder of the paper is structured as follows. Section 2 provides a brief analytical treatment of the safety-first model of insurance pricing and Dietz and Walker’s (2019) α -maxmin model of reserving under ambiguity, as well as how the risk of ruin is estimated using popular methods of model blending. Section 3 introduces and analyses the two case studies, and Section 4 provides a discussion.

2 Pricing insurance

2.1 Premium principle

An insurer that seeks to maximise its expected profits subject to a survival constraint would price premiums according to

$$p_c = L_c + y(Z_{f'} - Z_f), \tag{1}$$

where p_c is the premium on contract c , L_c is the expected loss, y is the cost of capital and Z_i represents the insurer’s capital reserves held against portfolio $i \in \{f, f'\}$. When the insurer takes on contract c , the insurer’s portfolio is $f' = f + c$. Hence the premium is comprised of the expected loss

on the contract, plus the cost of the additional capital required to ensure the portfolio-wide survival constraint is still met.

2.2 Expected losses and capital loads

Assuming a *single* probability distribution is appropriate to characterise the insurer’s losses on its portfolio, the expected loss L_c is precisely estimated, as are the capital reserves, which can be set so that they are just sufficient to cover the loss $(-)x$ at a pre-specified probability θ (e.g. 1/200 years or 0.005).⁷ For portfolio f this is

$$Z_f = \min \{x : P_f(-x) \leq \theta\}, \quad (2)$$

where $P_f(x) \equiv P_f(y : y < x)$, i.e. it is shorthand for the probability that portfolio f pays out any amount less than x , or equivalently that the loss is more than x .

The key question is what to do when multiple conflicting estimates of $P_f(x)$ are available and there is no basis for assuming one estimate is precisely correct. This might often be the case when insuring catastrophe risks such as hurricanes. The insurer may have at its disposal a set of estimates, one from each of the various catastrophe modelling companies.

One approach is to blend the various estimates into a single probability distribution and proceed exactly as above. There are several ways to blend models, but the principal methods are frequency and severity blending.⁸ Their workings will be described below.

Dietz and Walker (2019) propose an alternative approach. Where π is an estimate of the probability distribution of losses – call it a ‘model’ for short – the expected loss L_c is simply the expectation of the expected losses estimated by each π . On the other hand, in setting capital reserves the insurer may not be neutral towards the existence of ambiguity and this

⁷For the sake of consistency, we use the same notation as Dietz and Walker (2019). In that paper, x signifies the payout on a portfolio and hence $-x$ constitutes a loss.

⁸Like Calder et al. (2012) we take model blending to mean combining the outputs of different models, rather than combining the components of different models into a model that yields a single output, which can be referred to as ‘model fusion’.

implies not simply taking expectations. Thus, applying recent developments in economic theory, Dietz and Walker (2019) propose that capital reserves be set according to

$$Z_f = \min \left\{ x : \alpha \cdot \left[\max_{\pi \in \Pi} P_f^\pi(-x) \right] + (1 - \alpha) \cdot \left[\min_{\pi \in \Pi} P_f^\pi(-x) \right] \leq \theta \right\}, \quad (3)$$

where $\Pi \in \mathbb{Z}^+$ is the set of all models. The insurer computes Z_f by taking a weighted average of the highest and lowest estimates of the loss $(-x)$ at probability θ . The weight factor α captures the insurer's attitude to ambiguity.⁹

Dietz and Walker (2019) go on to show that, if one portfolio is more ambiguous than another,¹⁰ then an insurer holds more capital if and only if $\alpha > 0.5$. In other words, an insurer with $\alpha > (<)0.5$ is ambiguity-averse (-seeking) and charges higher (lower) premiums for contracts that increase the ambiguity of the portfolio. This comes about because an insurer with $\alpha > 0.5$ places more weight on the highest loss estimate, the worst case. In the limit of $\alpha = 1$, the insurer sets its capital reserves based exclusively on the worst case, which is analogous to (unweighted) maxmin decision rules that have been proposed as a means of making rational decisions under ambiguity/ignorance (e.g. Gilboa and Schmeidler, 1989; Hansen and Sargent, 2008).

⁹In Dietz and Walker (2019), this is denoted $\hat{\alpha}$ in order to differentiate it from the parallel concept in the underlying theory of decision-making under ambiguity (Ghirardato et al., 2004). However, doing so is unimportant here, so we avoid the extra notation.

¹⁰Formally, one portfolio f is more ambiguous than another portfolio g whenever any ambiguity-neutral insurer is indifferent between the two portfolios, any ambiguity-averse insurer prefers g to f , and any ambiguity-seeking insurer prefers f to g (Jewitt and Mukerji, 2017). Therefore this is a behavioural definition of ambiguity, depending only on revealed preferences, in keeping with the standard approach in economics.

2.3 Model blending

Frequency blending works by taking the weighted average of the probabilities estimated by each of the set of models Π of a given loss:

$$P_f(-x) = \sum_{\pi} \gamma_{\pi} P_f^{\pi}(-x), \quad (4)$$

where γ_{π} is the weight assigned to model π (see Figure 1 for a schematic representation). Clearly if each model weight in the set is equal to $1/\Pi$, then (4) is equivalent to computing the arithmetic mean of the probabilities at a given loss, which is a common and natural starting point.

Severity blending involves taking the weighted average of the losses estimated by the models at a given probability:

$$P_f(-x) = P_f \left[\sum_{\pi} \gamma_{\pi} \text{inv}P_f^{\pi}(-x^{\pi}) \mid P_f^j = P_f^k, \forall \pi = j, k \right]. \quad (5)$$

See Figure 1. If each model weight is equal to $1/\Pi$, then (5) is equivalent to computing the arithmetic mean of the losses at a given probability.

Although severity blending might appear to be the inverse of frequency blending, it is not. It is well known that the two techniques tend to produce different estimates of the composite loss distribution, even in relatively trivial examples where $\Pi = 2$ and $\gamma_1 = \gamma_2 = 0.5$ (see Calder et al., 2012, p18-21). Although severity blending involves applying the inverse of the loss distribution function, $\text{inv}P_f^{\pi}(-x)$, the weighted average loss at a given probability need not equal the weighted average probability at a given loss. Only in two cases will frequency and severity blending yield exactly the same estimate of $P_f(-x)$. One is the trivial case where the models agree exactly on the probability of a given loss, $P_f^j = P_f^k, \forall \pi = j, k$. The other is when the slopes of the loss distribution functions are equal, that is, for all pairs of models j and k , when $P_f^{j'}(-x^j) = P_f^{k'}(-x^k)$ over the interval $[-x^j, -x^k]$. See Appendix A.

FIGURE 1 ABOUT HERE

It has been argued that frequency blending is superior to severity blend-

ing. While severity blending is easy to perform, doing so breaks the link with the underlying event set, which makes it difficult to make comparisons, such as the accumulation of risk from a given peril across the insurer’s whole portfolio, or the comparison between losses gross and net of excess (Calder et al., 2012; Cook, 2011). Nevertheless both techniques are common and so we will evaluate both.

It is possible to envisage a situation in which the application of the α -maxmin reserving rule in (3) is mathematically equivalent to using frequency blending to mix models so that reserving rule (2) can be applied (or to using severity blending under the limited circumstances described above). This would be the case if, for example, there were two models and the ambiguity parameter α happened to coincide with the model weights $\{\gamma_\pi\}_{\pi=1}^2$. For instance, the insurer might be ambiguity-neutral ($\alpha = 0.5$) and consider each of the models to be equally likely to be the true model ($\gamma_1 = \gamma_2 = 0.5$). However, it is important to stress that α is a behavioural/preference parameter that is intended to capture the insurer’s attitude to ambiguity, whereas the model weights γ_π reflect the insurer’s beliefs about how likely each model is to be the true model. They are conceptually quite different.

3 Applications to hurricane risk

In this section, we estimate premiums in two case studies, both of which involve catastrophe modelling of hurricane risk and both of which are affected by ambiguity. On the one hand, we estimate premiums based on a single estimate of the probability of losses. This can be obtained either conditional on each model, by frequency blending or severity blending. On the other hand, we estimate premiums obtained when the capital holding is set according to the α -maxmin rule (3). In both case studies, the problem is simplified to that of a stand-alone, large-contract-cum-small-portfolio. Therefore we abstract from issues created by the interaction of the new contract with an insurer’s existing portfolio, as well as from other complications such as deductibles and limits. For these and other reasons that will become apparent, the case studies are not entirely realistic. Limitations to the availability of

data constitute a barrier to full realism. Nonetheless our hope is that the case studies are plausible enough to demonstrate the utility of the α -maxmin rule relative to model blending.

3.1 Florida property

The first case considers ambiguity affecting the insurance of a large portfolio of ca. 5 million residential buildings in Florida against wind damage from hurricanes. The data are taken from Ranger and Niehörster (2012). Ambiguity comes from a set of models of hurricane activity, i.e. the occurrence probabilities of catastrophic hurricane events are imprecisely known. The set of models is generated by taking a set of global climate models and using different models – both dynamical¹¹ and statistical – to downscale the boundary conditions from the global climate models to Atlantic hurricane activity. Our set comprises 15 models. The estimates of hurricane activity from these 15 models are then used to set the occurrence probabilities of hurricane events (e.g. the event rates) in a single, simplified hurricane risk model (from Risk Management Solutions Inc.), yielding 15 competing exceedance probability (EP) curves.¹² We interpret the loss forecasts provided by Ranger and Niehörster for the year 2020 as the complete set of estimates of present-day losses. This implies all the ambiguity comes from modelling the hazard, whereas in reality ambiguity could also stem from modelling vulnerability (and of course the available set of models may not capture the full range of possible hurricane frequencies; see Section 4).

The 15 EP curves are used to calculate average annual losses (AALs, equivalent to L_c in Equation (1)) and losses at specific return periods, and these enable us to estimate premiums if we abstract from the pre-existing portfolio, i.e. we assume $Z_f = 0$ so that $Z_{f'} = Z_c$. Figure 2 displays the premium conditional on each hurricane model. It also displays ambiguity-adjusted premiums, with the insurer’s capital reserves set according to the α -

¹¹That is, models that use basic physical principles to calculate changes in climatic features.

¹²Ranger and Niehörster (2012) provide estimates assuming (i) status quo property vulnerability, and (ii) lower vulnerability based on all properties meeting 2004 Florida building codes. We use the status quo assumption.

maxmin rule (3), as well as premiums based on using frequency and severity blending of the models to obtain the AAL and the insurer’s reserves via Equation (2).¹³ Each of the 15 model EP curves is given equal weight when carrying out frequency and severity blending, i.e. $\gamma_\pi = 1/H, \forall \pi$. Clearly other model weighting schemes are possible and historical data could be used for that purpose (e.g. from HURDAT). Equal weighting is consistent with the principle of insufficient reason and the view of Ranger and Niehörster was that there was no straightforward way of establishing relative confidence in each model. In all cases, the insurer is required to hold capital sufficient to cover a 1/200-year event and the cost of capital is assumed to be 10%.

FIGURE 2 ABOUT HERE

The range of premiums estimated is from \$8.8 billion under the UKMO-HADCM3 model (a global climate model dynamically downscaled) to \$16.3bn under the MDR-SST model (a statistical model predicting hurricane formation based on the sea surface temperature in the Atlantic Main Development Region). An ambiguity-neutral insurer with $\alpha = 0.5$ requires a premium of \$11.3bn. This is an important benchmark, since an ambiguity-averse (-seeking) insurer requires a strictly higher (lower) premium. The premium of \$11.3bn is comprised of an AAL of \$3.8bn and a capital load of \$7.5bn, thus the premium is three times the expected loss (expectations taken over all 15 models). A maximally ambiguity-averse insurer with $\alpha = 1$ sets its premium equal to the expected AAL plus the capital required to absorb the 1/200-year loss in the most pessimistic model, MDR-SST. The resulting premium is \$13.6bn, so the ambiguity load is \$2.2bn or 20%, and the price multiple of the expected loss rises to 3.5.¹⁴ Note that this premium is lower than the premium conditional on MDR-SST itself, because the expected AAL over all 15 models is lower than the AAL predicted by MDR-SST.

¹³As the underlying model event sets were unavailable to us, frequency blending was performed by directly computing (4) at each loss using the models’ EP curves.

¹⁴These premiums/multiples may be compared with the empirical estimates of Lane and Mahul (2008), according to which a representative multiple for US wind based on the market for insurance linked securities is 3.3, averaging over the insurance cycle. However, since we abstract from the pre-existing portfolio, one should not read too much into the comparison.

One of the key results contained in Figure 2 is that frequency and severity blending produce a premium lower than Equation (3) when $\alpha = 0.5$. This implies frequency and severity blending are consistent with an insurer who is *ambiguity-seeking* in this case. The premium is particularly low under frequency blending; it is lower than the premium conditional on 12/15 models. The shortfall between frequency/severity blending and the ambiguity-neutral premium is \$1.8bn and \$0.8bn respectively. Inspecting the figure, it is clear that a contributing factor to this result is that the most pessimistic model, MDR-SST, which is given a weight of 0.5 by an ambiguity-neutral insurer in setting its capital reserves, is an outlier. It is therefore tempting to dismiss the relevance of the result to cases where the most pessimistic model is not an outlier. However, note that if MDR-SST is (arbitrarily) removed from the set of models, while the resulting premium under severity blending is now just above the ambiguity-neutral level, the premium under frequency blending remains below it.¹⁵

Figure 3 shows that the frequency-blended EP curve gives lower losses at lower return periods than the severity-blended EP curve, and *vice versa* at higher return periods. Applying the theory set out in Appendix A, this is driven by the slope of the more pessimistic EP curves being initially higher, resulting in an evident dispersion of loss estimates, which ceases to be the case as the return period goes above 250 years. It is these lower losses at lower return periods that are responsible for the particularly low premiums under frequency blending.

FIGURE 3 ABOUT HERE

In Figure 4, we repeat the analysis, but instead require insurers to hold sufficient capital to survive a 1/500-year loss. All premiums and multiples are higher, but the ambiguity load for ambiguity-averse insurers is uniformly slightly smaller. For $\alpha = 1$, the multiple of the price above the expected loss is 4.5, but the ambiguity load is 16%. This reflects the fact that the models are marginally less dispersed in their estimates of the 1/500-year event than

¹⁵An ambiguity-neutral insurer with $\alpha = 0.5$ requires a premium of \$10.07bn when MDR-SST is excluded. Severity blending yields a corresponding premium of \$10.11bn, while frequency blending yields a comparable premium of \$9.3bn.

they are of the 1/200-year event.¹⁶ Frequency and severity blending still give rise to premiums consistent with an ambiguity-seeking insurer. Notice that a maximally ambiguity-seeking insurer with $\alpha = 0$ charges a premium below that conditional on the most optimistic model, UKMO-HADCM3. This is explained by the fact that, although UKMO-HADCM3 forecasts the lowest AAL, it does not forecast the lowest 1/500-year losses (the model WND SHR&MDR-SST does¹⁷).

FIGURE 4 ABOUT HERE

3.2 Dominica property

The second case considers ambiguity affecting the insurance of a portfolio of 3,000 low-income residential properties on the Caribbean island of Dominica against hurricane wind damage. The data come from a study by RMS et al. (2018). In this case there is a single hazard model and ambiguity comes from a set of six vulnerability functions, based on different assumptions about the construction quality of buildings in the portfolio.

As in case 1, the six EP curves are used to calculate AALs, losses at specific, high return periods against which capital reserves are set, and then premiums, again assuming $Z_{f'} = Z_c$. Figure 5 displays premiums where capital reserves are set against the 1/200-year event and the cost of capital is 10%. The six models are assigned equal weight for model blending. The range of premiums is from \$469,580 under the lowest vulnerability function to about \$1.4 million under the highest. An ambiguity-neutral insurer with $\alpha = 0.5$ requires a premium of \$934,322, which is five times the expected loss of \$184,516. The ambiguity-neutral premium is higher than the premiums conditional on 5/6 models, reflecting how much higher losses are under the highest vulnerability function (no protection) than the other five. This also explains why the multiple is as much as five. The maximum ambiguity load

¹⁶As two indications of this, the respective coefficients of variation are 0.10 and 0.16, while the respective adjusted Fisher-Pearson coefficients of skewness are 0.5 and 1.8.

¹⁷WNSHR&MDR-SST is a statistical model, which predicts hurricane formation based on the local vertical wind shear and the sea surface temperature in the Atlantic Main Development Region.

is \$397,769 or 43%, defined as before as the difference in premium between $\alpha = 0.5$ and $\alpha = 1$. In this case we again find that frequency and severity blending give rise to premiums that are below the premium required by the ambiguity-neutral insurer, thus being representative of ambiguity-seeking behaviour by the insurer.

FIGURE 5 ABOUT HERE

The implications of model blending for an insurer's attitude to ambiguity depend on the model weighting scheme. To demonstrate this, here we implement a purely illustrative scheme, according to which the model using the highest vulnerability function (no protection) is over-weighted at $\gamma_{\text{no}} = 0.4$, the three models using vulnerability functions that assume partial roof reinforcement are given a probability of 0.15, and the two models using the lower vulnerability functions (full retrofit) are assigned $\gamma_{\pi} = 0.075$. Figure 6 shows that in this case severity blending is consistent with insurer ambiguity aversion, but frequency blending still gives rise to a premium below the ambiguity-neutral premium, i.e. it is inconsistent with ambiguity aversion.

FIGURE 6 ABOUT HERE

This result suggests a relationship between the weight placed on the model using the highest vulnerability function and the premium required under model blending. Figure 7 traces that relationship, where γ_{no} is the probability assigned to no protection and the weights of the remaining five models are kept equal to $(1 - \gamma_{\text{no}}) / 5$. Figure 7 shows that the no protection model needs to be assigned a weight of $\gamma_{\text{no}} > 0.28$ if severity blending is to be consistent with ambiguity aversion, while it needs to be assigned a weight of $\gamma_{\text{no}} > 0.69$ if frequency blending is to be consistent with ambiguity aversion.

FIGURE 7 ABOUT HERE

4 Discussion

The two case studies in this paper have demonstrated the application of a premium principle, which is based on expected profit maximisation, com-

bined with a survival constraint that is based on an α -maxmin rule for mixing the least and most pessimistic estimates of the risk of ruin. The analysis has shown that the α -maxmin rule lends itself quite naturally to a setting in which there is a finite number of competing catastrophe models. The ambiguity load is up to 20% in the case of Florida property and 43% in the case of Dominica property for a cost of capital of 10% and a risk of ruin of 1/200, with the difference being explained principally by the dispersion in model estimates of the 1/200-year loss. Since we abstract from any pre-existing insurance portfolio and therefore assume away the possibility of any risk pooling at the portfolio level, these estimates should not be taken too literally. The main point is to show how the ambiguity load can be linked explicitly to the insurer's attitude to ambiguity, α .

This naturally begs the question of what an insurer's value of α should be and how it might be calibrated. This is not straightforward to answer at this stage. Since the notion of applying the α -maxmin rule to insurance pricing is new, no directly relevant empirical evidence on insurers' α exists.¹⁸ Insofar as the insurer wishes to set premiums consistent with aversion to ambiguity, it is nonetheless clear that $0.5 < \alpha \leq 1$ and this provides a good place to start. Within this range, our suggestion is to take a reflexive approach, whereby the insurer considers the ambiguity load (and related information such as the insurance multiple) that is implied by various values of α , and uses the results to determine what the insurer thinks its α should be. There is a certain parallel here with how Ellsberg (1961) and others (e.g. Gilboa et al., 2009) have sought to justify ambiguity aversion as a feature of rational behaviour in the round: namely confront subjects with their behaviour in Ellsberg-type experiments and get them to reflect on whether this behaviour appears rational to them.

Another difficult question posed by the analysis is; how would insurers know that the set of models they are looking at, \mathcal{I} , is complete? It is nat-

¹⁸The experimental/survey studies discussed in Section 1 do provide evidence on the size of the ambiguity load within the context of the scenarios used, but the α -maxmin rule links the ambiguity load with the insurer's capital, which is not considered in these studies.

urally difficult to define the things one doesn't know. There is an emerging literature in economic theory on unawareness, particularly awareness of unawareness, in which rational decision-makers are modelled as being aware that their perception of the decision problem could in some way be incomplete, i.e. principally that their subjective state space could be incomplete (Walker and Dietz, 2011; Lehrer and Teper, 2014; Alon, 2015; Grant and Quiggin, 2015; Karni and Vierø, 2017; Kochov, 2018). All of these theories allow for precautionary behaviour in such situations (some require it). It may in future be possible to apply some of these frameworks to insurance (e.g. the framework of Alon, 2015), although at present they start from the position of an expected utility maximiser, so it is unclear how they can be reconciled with a framework in which there is some degree of well-defined ambiguity to start with. If the insurer suspects that \mathcal{I} is incomplete, what are the implications for pricing? One practical implication is that, insofar as the insurer wishes to exercise further caution in the face of awareness of its unawareness, the continued use of an *ad hoc* premium multiplier could be justified, if the insurer suspects losses could exceed those in the most pessimistic model.

In practice, it may be infeasible to include the full set of available catastrophe models when estimating reserves and premiums, due for instance to cost, or a lack of in-house expertise to run them. The question then arising is whether a particular subset of models would suffice. Using data from the Florida property case, Figure 8 shows the premium estimated by applying the α -maxmin rule to the full set of models (left-hand side) and compares this with the premium estimated when one of the 15 hurricane models is omitted (i.e. the premium is estimated over $\mathcal{I} - 1$ models). This is repeated for each of the 15 models. In addition, on the right-hand side is the premium estimated using just the least and most pessimistic models (with respect to the 1/200-year loss). The figure illustrates two points. First, given the way in which the α -maxmin rule takes a convex combination of the least and most pessimistic models to estimate the capital load, the premium is naturally relatively insensitive to removing any model except these two. This suggests it is important for insurers to acquire the least and most pessimistic

models, insofar as this is known. Second, just using the least and most pessimistic models does not, however, provide a good approximation of the full set, at least in this case. This is because combining just these two models significantly overestimates the AAL. This in turn stems from the fact that, while WNSHR&MDR-SST has the lowest 1/200 loss, it actually has quite a high AAL.

FIGURE 8 ABOUT HERE

A more straightforward implication of the analysis is that frequency and severity blending can lead to premiums that are inconsistent with insurer ambiguity aversion. This follows straightforwardly from the observation that the premium estimated by these two blending methods is often lower than that estimated by the α -maxmin rule when $\alpha = 0.5$. For this result, we need only interpret premium pricing under the α -maxmin rule as a representation, rather than a recommendation. Clearly whether there is an inconsistency is context-dependent. All we have sought to do is provide plausible examples.

If the insurer would subsequently apply an *ad hoc* adjustment to the premium estimated via model blending, the premium could still be consistent with ambiguity aversion, of course. We can use the results above to illustrate how large this adjustment would need to be. Taking the leading examples presented in Figures 2 and 5, the multiplier on the premium estimated via frequency blending would need to be at least 1.15 for the Florida property portfolio and 1.3 for the Dominica property portfolio. If severity blending is used, then the resulting multipliers would need to be at least 1.08 for the Florida property portfolio and 1.09 for the Dominica property portfolio. Nonetheless in our view the use of *ad hoc* multipliers is better limited to conservatism in the face of unspecified contingencies, as discussed above.

References

Alary, D., C. Gollier, and N. Treich (2013). The effect of ambiguity aversion on insurance and self-protection. *Economic Journal* 123(573), 1188–1202.

- Alon, S. (2015). Worst-case expected utility. *Journal of Mathematical Economics* 60, 43–48.
- Berger, L. (2016). The impact of ambiguity and prudence on prevention decisions. *Theory and Decision* 80(3), 389–409.
- Berger, L. A. and H. Kunreuther (1994). Safety first and ambiguity. *Journal of Actuarial Practice* 2(2), 273–291.
- Cabantous, L. (2007). Ambiguity aversion in the field of insurance: insurers’ attitude to imprecise and conflicting probability estimates. *Theory and Decision* 62(3), 219–240.
- Cabantous, L., D. Hilton, H. Kunreuther, and E. Michel-Kerjan (2011). Is imprecise knowledge better than conflicting expertise? evidence from insurers’ decisions in the united states. *Journal of Risk and Uncertainty* 42(3), 211–232.
- Calder, A., A. Couper, and J. Lo (2012). Catastrophe Model Blending: Techniques and Governance. Technical report, Institute and Faculty of Actuaries.
- Cook, I. (2011). Using multiple catastrophe models. Technical report, Institute and Faculty of Actuaries.
- Dietz, S. and O. Walker (2019). Ambiguity and insurance: capital requirements and premiums. *Journal of Risk and Insurance* 86(1), 213–235.
- Einhorn, H. and R. Hogarth (1985). Ambiguity and uncertainty in probabilistic inference. *Psychological Review* 92(4), 433–461.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics* 75(4), 643–669.
- Ghirardato, P., F. Maccheroni, and M. Marinacci (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118(2), 133–173.

- Gilboa, I., A. Postlewaite, and D. Schmeidler (2009). Is it always rational to satisfy savage's axioms? *Economics and Philosophy* 25(3), 285–296.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Gollier, C. (2014). Optimal insurance design of ambiguous risks. *Economic Theory* 57(3), 555–576.
- Grant, S. and J. Quiggin (2015). A preference model for choice subject to surprise. *Theory and Decision* 79(2), 167–180.
- Hansen, L. and T. Sargent (2008). *Robustness*. Princeton University Press.
- Hogarth, R. M. and H. Kunreuther (1989). Risk, ambiguity, and insurance. *Journal of Risk and Uncertainty* 2(1), 5–35.
- Hogarth, R. M. and H. Kunreuther (1992). Pricing insurance and warranties: ambiguity and correlated risks. *Geneva Papers on Risk and Insurance Theory* 17(1), 35–60.
- Jewitt, I. and S. Mukerji (2017). Ordering ambiguous acts. *Journal of Economic Theory* 171, 213–267.
- Karni, E. and M.-L. Vierø (2017). Awareness of unawareness: a theory of decision making in the face of ignorance. *Journal of Economic Theory* 168, 301–328.
- Kochov, A. (2018). A behavioral definition of unforeseen contingencies. *Journal of Economic Theory* 175, 265–290.
- Kunreuther, H., R. Hogarth, and J. Meszaros (1993). Insurer ambiguity and market failure. *Journal of Risk and Uncertainty* 7(1), 71–87.
- Kunreuther, H., J. Meszaros, R. Hogarth, and M. Spranca (1995). Ambiguity and underwriter decision processes. *Journal of Economic Behavior & Organization* 26(3), 337–352.

- Kunreuther, H. and E. Michel-Kerjan (2009). *At War with the Weather: Managing Large-Scale Risks in a New Era of Catastrophes*. Cambridge, MA: MIT Press.
- Landsea, C. W., N. Nicholls, W. M. Gray, and L. A. Avila (1996). Downward trends in the frequency of intense at Atlantic Hurricanes during the past five decades. *Geophysical Research Letters* 23(13), 1697–1700.
- Lane, M. and O. Mahul (2008). Catastrophe risk pricing: an empirical analysis. Technical report, World Bank Policy Research Working Paper 4765.
- Lehrer, E. and R. Teper (2014). Extension rules or what would the sage do? *American Economic Journal: Microeconomics* 6(1), 5–22.
- Machina, M. J. and M. Siniscalchi (2014). Ambiguity and ambiguity aversion. In *Handbook of the Economics of Risk and Uncertainty*, Volume 1, pp. 729–807. Elsevier.
- Mao, M. (2014). Catastrophe model blending. In *CAS Ratemaking and Product Management Seminar*.
- Pichler, A. (2014). Insurance pricing under ambiguity. *European Actuarial Journal* 4(2), 335–364.
- Ranger, N. and F. Niehörster (2012). Deep uncertainty in long-term hurricane risk: scenario generation and implications for future climate experiments. *Global Environmental Change* 22(3), 703–712.
- RMS, Vivid Economics, re:focus partners, DfID, and Lloyd’s (2018). Financial instruments for resilient infrastructure. Technical report.
- Snow, A. (2011). Ambiguity aversion and the propensities for self-insurance and self-protection. *Journal of Risk and Uncertainty* 42(1), 27–43.
- Stone, J. (1973). A theory of capacity and the insurance of catastrophe risks (part i). *Journal of Risk and Insurance*, 231–243.

Walker, O. and S. Dietz (2011). A representation result for choice under conscious unawareness.

Figure 1: Schematic representation of frequency and severity blending of exceedance probability curves




Figure 2: Premiums for a portfolio of Florida residential property under event rates derived from different hazard models




Figure 3: Exceedence probability curves under 15 models of hurricane formation and based on frequency and severity blending (vertical axis is truncated at \$120bn for ease of inspection)




Figure 4: Premiums for a portfolio of Florida residential property, where insurers hold capital to cover 1/500-year loss




Figure 5: Premiums for a portfolio of Dominican residential property under different vulnerability models




Figure 6: Premiums for Dominican residential property with an illustrative scheme of unequal model weights




Figure 7: Premiums for Dominican property as a function of weight assigned to highest loss model, no protection





Figure 8: Premiums for a portfolio of Florida residential property under different model subsets



A When do frequency and severity blending give the same result?

In order for frequency and severity blending to yield the same estimate of $P_f(-x)$, from Equations (4) and (5) it must be that

$$\sum_{\pi} \gamma_{\pi} P_f^{\pi}(-x) = P_f \left[\sum_{\pi} \gamma_{\pi} \text{inv} P_f^{\pi}(-x^{\pi}) \mid P_f^j = P_f^k, \forall \pi = j, k \right].$$

Using Jensen's inequality, a necessary and sufficient condition for this is that $P_f^{j'}(-x^j) = P_f^{k'}(-x^k)$ for all j, k over the interval $[-x^j, -x^k]$, i.e. all model loss distribution functions have the same slope, or in other words the reduction in probability is the same for a given increase in loss. This is trivially the case when $P_f^j(-x) = P_f^k(-x)$ for all j, k , i.e. all models agree and there is no ambiguity, but may occasionally be the case in other situations too.


Take any pair of models and suppose the more pessimistic model is j and the more optimistic model is k , in the sense that j estimates a higher probability of a given loss. Then Jensen's inequality also implies

$$\sum_{\pi} \gamma_{\pi} P_f^{\pi}(-x) \geq (\leq) P_f \left[\sum_{\pi} \gamma_{\pi} \text{inv} P_f^{\pi}(-x^{\pi}) \mid P_f^j = P_f^k, \forall \pi = j, k \right]$$

if and only if $P_f^{j'}(-x^j) \geq (\leq) P_f^{k'}(-x^k)$ for all j, k over the interval $[-x^j, -x^k]$. This says – in strong form – that frequency blending yields at least as high an estimate of the probability of loss $-x$ if and only if the slope of the more pessimistic loss distribution function j is higher than its counterpart k , for all j, k , and *vice versa*. Figure 9 provides a schematic representation. If this is true, then given slope of the more pessimistic function is always higher, we will observe an increasing dispersion of model estimates over the relevant range of losses.

Sometimes, as in Section 3, it is more intuitive to perform the inverse mapping of return periods to losses. Then the logic is the inverse too:

Figure 9: Schematic representation of Jensen's inequality as applied to model blending



severity blending yields a higher estimate of the probability of loss $-x$ if and only if the slope of the more pessimistic *inverse* loss distribution function j is higher than its counterpart k , for all j, k , and *vice versa*. This intuition is used to explain what we see in Figure 3.