LSE Research Online

Simon Dietz and Oliver Walker

Ambiguity and insurance: capital requirements and premiums

Article (Accepted version)
(Refereed)

This version available at: http://eprints.lse.ac.uk/68469/
Available in LSE Research Online: December 2016

http://eprints.lse.ac.uk

Ambiguity and insurance: capital requirements and premiums

Simon Dietz[*][†] and Oliver Walker[†][‡]

November 21, 2016

Abstract

Many insurance contracts are contingent on events such as hurricanes, terrorist attacks or political upheavals, whose probabilities are ambiguous. This paper offers a theory to underpin the large body of empirical evidence showing that higher premiums are charged under ambiguity. We model a (re)insurer who maximises profit subject to a survival constraint that is sensitive to the range of estimates of the probability of ruin, as well as the insurer's attitude towards this ambiguity. We characterise when one book of insurance is more ambiguous than another and general circumstances in which a more ambiguous book requires at least as large a capital holding. We subsequently derive several explicit formulae for the price of insurance contracts under ambiguity, each of which identifies the extra ambiguity load.

[*]ESRC Centre for Climate Change Economics and Policy and Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science.
[†]Department of Geography and Environment, London School of Economics and Political Science.
[‡]Vivid Economics Ltd, London.

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

1 Introduction

Many (re)insurance contracts are contingent on events such as hurricanes, terrorist attacks or political upheavals whose probabilities are not known with precision. Such contracts are said to be subject to "ambiguity". There may be several reasons why contracts are subject to ambiguity, including a lack of historical, observational data, and the existence of competing theories, proffered by competing experts and formalised in competing forecasting models, of the causal processes governing events that determine their value. For example, ambiguity is a salient feature in the insurance of catastrophe risks such as hurricane-wind damage to property in the southeastern United States. Here, historical data on the most intense hurricanes are limited, and there are competing models of hurricane formation (Bender et al., 2010; Knutson et al., 2008; Ranger and Niehoerster, 2012). This ambiguity is increased by the potential role of climate change in altering the frequency, intensity, geographical incidence and other features of hurricanes.

There is by now a body of evidence to show that, faced with offering a contract under ambiguity, insurers increase their premiums, limit coverage, or are unwilling to provide insurance at all. Much of the academic evidence is survey-based: actuaries and underwriters from insurance and reinsurance companies are asked to quote prices for hypothetical contracts in which the probabilities of loss are alternatively known or unknown (Cabantous, 2007; Cabantous et al., 2011; Hogarth and Kunreuther, 1989, 1992; Kunreuther et al., 1993, 1995; Kunreuther and Michel-Kerjan, 2009). Their responses reveal that prices for contracts under ambiguity exceed prices for contracts without ambiguity and with equivalent expected losses, which is consistent with ambiguity aversion[1] and thus in line with a much larger body of evidence on decision-making, starting with Ellsberg's classic thought experiments on choices over ambiguous and unambiguous lotteries (Ellsberg, 1961). In the industry, one can find guidance that insurers should increase their 'prudential margins' (i.e. capital holdings) under ambiguity (e.g. Barlow et al., 1993) and below we explain how this leads to higher premiums.

Yet, despite the evidence, there is seemingly little theoretical work that can explain or formally motivate these ambiguity loadings. In this paper we seek to fill this hole by offering a formal analysis of the connection between, on the one hand, ambiguous information about the performance of a book of insurance and, on the other hand, the premium charged for a new contract. We do so via the capital held against the book: our starting point is a well-known model of the price of insurance, according to which the objective is to maximise expected profits subject to a survival constraint (thus in the tradition of Stone, 1973), which is imposed by managerial or regulatory fiat

---

[1] We give formal definitions of ambiguity, ambiguity aversion and related concepts later.

2

out of concern for ensuring solvency or avoiding a downgrading of credit. An example of such a constraint, imposed by regulation, is the European Union's new Solvency II Directive (where it is called a Solvency Capital Requirement). Our twist is that the capital held is sensitive to the range of estimates of the probability of ruin and to the insurer's attitude towards ambiguity in this sense.

Based on recent contributions to the theory of decision-making under ambiguity, we characterise circumstances in which one book of insurance is "more ambiguous" than another, and establish general conditions under which more ambiguous books entail higher capital holdings under our capital-setting rule. We then use the rule to derive pricing formulae for ambiguous contracts in a way that isolates the additional *ambiguity load*, distinct from the more familiar risk load. We examine the properties of the ambiguity load under different assumptions about the insurer's information: it is shown to depend on the ambiguity of the contract being priced, as well as the insurer's ambiguity aversion. It also depends on the relationship between the ambiguity of the new contract and the ambiguity of the pre-existing book, while under some circumstances it can interact with the coventional risk load. We hope that these pricing formulae, or further extensions and refinements of them, may prove practically useful in the industry: one of the consequences of the lack of existing theory is that the practice of loading contract prices under ambiguity does not appear to have been codified and may often use back-of-the-envelope calculations and heuristics (Hogarth and Kunreuther, 1992).

Our paper is a complement to recent work on how ambiguity, and ambiguity aversion, on the part of would-be policyholders affects the characteristics of optimal insurance contracts (Alary et al., 2013; Gollier, 2014). In this work, the insurer is taken to be ambiguity-neutral, whereas our insurer is ambiguity-averse. Our paper is also related to recent work on ambiguity aversion and robust control that has taken a similar approach, but applied it to different problems. Notable examples include Garlappi et al. (2007) on portfolio selection, and Zhu (2011) on catastrophe-risk securities. Finally, our paper offers an alternative approach to previous work in the literature on insurance that has also considered ambiguity under the auspices of 'model uncertainty' (e.g. Cairns, 2000). This work also features ambiguity-neutral insurers, because it is assumed that they reduce compound lotteries à la probabilistic sophistication (Machina and Schmeidler, 1992; Epstein, 1999).

The rest of the paper is organised as follows. Section 2 presents the decision problem formally. Section 3 considers the relationship between how ambiguous a book of insurance is and how much capital the insurer must hold, drawing on elements of Jewitt and Mukerji's (2012) characterisation of the "more ambiguous" relation. Section 4 then derives explicit pricing formulae

for insurance contracts under ambiguity. Finally, Section 5 concludes with a discussion of the descriptive and normative appeal of our capital-setting rule, and some interpretation of our results.

2 The insurer's decision problem

We take the point of view of an insurer who faces uncertainty over the performance of its book and wishes to maximise expected profits subject to a survival constraint. The classic treatment of this problem characterises the insurer's uncertainty using a single probability measure over the space of events determining the book's return (e.g. Kreps, 1990). Under this account the insurer may control the likelihood of insolvency/ruin by choosing a capital holding, since the likelihood of ruin is then simply the probability that the book's losses are not covered by the capital.

As we explained in the Introduction, there are, however, important cases where the insurer's information does not take the form of a single probability measure over a space of relevant scenarios. In such cases, it may entertain a multiplicity of possible measures over the space of payoff-relevant events and not be certain which of them "correctly" quantifies the uncertainty it faces. Such an insurer is said to face *ambiguity*.

We model this kind of insurer's information as follows. There is a metric space, S, known as the *state space*, that consists of all of the possible states of the world that are relevant to the performance of an insurance book, with the Borel σ-algebra on S denoted $\mathcal{B}$. A *book* is then a $\mathcal{B}$-measurable mapping from S to $\mathbb{R}$. We denote the full set of books by $\mathcal{F}$ and, where $f \in \mathcal{F}$, interpret $f(s) = x$ as the statement that if s turns out to be the true state of the world, book f will return the monetary quantity x. A book is thus identical to a "Savage act" (in the sense of Savage, 1954) and, in terms of the insurer's capital setting and pricing, both the pre-existing insurance portfolio and the new contract to be potentially added to this portfolio may be regarded as books.

In the classic account of this problem, the insurer is assumed to possess a single probability measure on $\mathcal{B}$, representing its information about payoff-relevant events. We, however, wish to allow for cases where the insurer faces ambiguity and therefore endow it with a *set* of measures on $\mathcal{B}$, Π, encompassing all probability measures it believes might characterise its uncertainty correctly. We refer to Π as the set of *models*, and require Π to be compact and convex. Where the insurer's book depends, for example, on weather events, Π may consist of a set of seasonal forecasts, one of which is assumed to be correct insofar as it accurately measures the likelihood of any member of $\mathcal{B}$. Where $\mathcal{B}_\Pi$ is a Borel σ-algebra on Π, let ν be the probability measure on $\mathcal{B}_\Pi$ to represent the insurer's beliefs about which of the models in Π is

correct. We require $\text{supp}(\nu) = \Pi$.

Using $\mathcal{B}_{\mathbb{R}}$ for the Borel σ-algebra on $\mathbb{R}$, for any $f \in \mathcal{F}$ we can define the probability measure P_f on $\mathcal{B}_{\mathbb{R}}$ as follows:

$$P_f(E) = \int_{\Pi} \pi \left(f^{-1}(E) \right) d\nu$$

for any $E \in \mathcal{B}_{\mathbb{R}}$. In words, given the insurer's beliefs about Π, $P_f(E)$ gives the probability the insurer places on its book paying out some amount in E. Throughout this paper, we adopt the convention of using $P_f(y)$ for $P_f(\{x : x < y\})$: the probability, given beliefs ν over the measures Π, that f pays out less than y.

Let us take as our starting point a familiar model of insurance pricing based on maximising expected profit subject to a survival constraint, in the tradition of Stone (1973).[2] In this model, there is an insurer who, given any book $f \in \mathcal{F}$, sets its capital holding, Z_f, as follows:

$$Z_f = \min\{x : P_f(-x) \le \theta\} \tag{1}$$

That is, Z_f is the smallest holding such that the probability of losses exceeding it is no more than some benchmark level θ (we take it for granted here and throughout the paper that (1) well defines Z_f). Given how we define $P_f(.)$, one can alternatively think of x as the Value at Risk of book f with respect to the "confidence level" $1 - \theta$. The requirement that the insurer holds Z_f may be interpreted as a managerial or regulatory constraint with the magnitude of θ representing the conservatism of the regime responsible for it. The insurer thus focuses on the single probability – as measured by P_f – of its book paying out less than its capital holding.

We extend this framework to allow the capital holding to depend on both the range of models Π and the insurer's attitude to ambiguity about the risk of ruin. Specifically, our insurer sets Z_f according to

$$Z_f = \min \left\{ x : \hat{\alpha} \cdot \left[\max_{\pi \in \Pi} P_f^{\pi}(-x) \right] + (1 - \hat{\alpha}) \cdot \left[\min_{\pi \in \Pi} P_f^{\pi}(-x) \right] \le \theta \right\} \tag{2}$$

where $\hat{\alpha} \in [0, 1]$, and the measure P_f^{π} on $\mathcal{B}_{\mathbb{R}}$ is defined as $P_f^{\pi}(E) = \pi \left(f^{-1}(E) \right)$. In contrast to (1), (2) requires the insurer to consider the dispersion in Π. [3]

---

[2]Ignoring, however, his stability constraint on the volatility of the ratio of losses to expenses.

[3]Notice that (2) does not require a probability measure ν on $\mathcal{B}_{\Pi}$, rather it suffices to know the models that yield respectively the maximum and minimum probability of the book paying out less than the insurer's capital holding. However, to say in Section 3 that one book is more or less ambiguous than another, we do require such a probability measure. We also require it to explicitly identify the ambiguity load in the premium price, which is the purpose of Section 4.

The weight factor $\hat{\alpha}$ plays a role akin to the ambiguity attitude parameter α in the well known α-maxmin expected utility (α-MEU) representation of choice under ambiguity of Ghirardato et al. (2004). In particular, we show in the next section that a more ambiguous book, which can be defined in terms of the preferences under ambiguity of an α-MEU decision-maker *inter alia*, incurs a higher capital holding Z_f if and only if the weight factor $\hat{\alpha} \geq 0.5$. Therefore $\hat{\alpha}$ indexes the insurer's aversion to ambiguity about the risk of ruin. Notice that if $\hat{\alpha} = 1$ then (2) simplifies to $Z_f = \min\{x : \max_{\pi \in \Pi} P_f^\pi(-x) \leq \theta\}$, which encodes a concern for robustness in a way analogous to the decision rule in Hansen and Sargent (2008), which itself has been shown to be equivalent to Gilboa and Schmeidler's (1989) axiomatically founded "maxmin" expected utility representation (Hansen and Sargent, 2001).

The capital holding affects the premium charged on a new contract, added to the existing portfolio, in the following way. An insurer endowed with book f who agrees to an additional contract c, itself a book, ends up with book $f' = f + c$. We define the addition operation over $\mathcal{F}$ pointwise – that is, for $f, c \in \mathcal{F}$, $f + c = f'$ where $f'(s) = f(s) + c(s)$ for all s – and note that $\mathcal{F}$ is closed under addition – i.e. if $f, c \in \mathcal{F}$, $f + c \in \mathcal{F}$. As a result of signing c, it needs to increase its capital holding by $Z_{f'} - Z_f$, and if c is competitively priced, then the underwriter's expected profit from the contract cannot exceed the opportunity cost of this incremental capital holding. Thus, using y for the opportunity cost of capital, if c is competitively priced it must be that:

$$\mu_c = y\left(Z_{f'} - Z_f\right) \tag{3}$$

The expected return on c, μ_c, is equal to the price the insurer charges the counterparty to c, p_c, less the expected loss on c (including administrative costs), L_c, and we say an insurer is *competitive*[4] whenever it always sets p_c such that μ_c satisfies (3):

$$p_c = L_c + y\left(Z_{f'} - Z_f\right) \tag{4}$$

The purpose of Section 4 is to expand the pricing formula (4) so that the *ambiguity load* – the uplift on the premium due to ambiguity about the new contract and the existing book – can be isolated under various distributional assumptions.

3   Ambiguity and the capital holding

Jewitt and Mukerji (2012) provide various choice-based accounts of what it is

---

[4]Note this our approach does not require the assumption of perfect competition. y may be interpreted as a managerial target rate of return rather than opportunity cost, in which case it could be consistent with a monopolistic or oligopolistic insurance industry.

for one act – book of insurance in our context – to be "more ambiguous" than another. We focus on one of these accounts, according to which book f is more ambiguous than g whenever any ambiguity-neutral agent is indifferent between the two books, any ambiguity-averse agent prefers g to f, and any ambiguity-seeking agent prefers f to g. Note that under this definition what it takes for f to be more ambiguous than g depends on what it means for an agent to be ambiguity-averse, -seeking, or -neutral, in keeping with revealed-preference traditions.

To characterise ambiguity attitude we primarily use Ghirardato, Maccheroni and Marinacci's (2004, GMM) axiomatically based α-MEU representation of choice under ambiguity, according to which preferences, given by the relation $\succeq$ over $\mathcal{F}$, are such that, for any bounded $f, g \in \mathcal{F}$:

$$
\begin{aligned}
f \succeq g \iff \quad & \alpha \cdot \min_{\pi \in \Pi} \int_S u\left(f(s)\right) d\pi + (1-\alpha) \cdot \max_{\pi \in \Pi} \int_S u\left(f(s)\right) d\pi \geq \\
& \alpha \cdot \min_{\pi \in \Pi} \int_S u\left(g(s)\right) d\pi + (1-\alpha) \cdot \max_{\pi \in \Pi} \int_S u\left(g(s)\right) d\pi
\end{aligned} \tag{5}
$$

where $\alpha \in [0, 1]$ is an index of ambiguity attitude, u is a continuous, non-decreasing function representing risk attitude in the usual way, and Π, which is compact and convex, represents the beliefs revealed by $\succeq$. Where $\succeq_A$ and $\succeq_B$ belong to the class of α-MEU preferences with common beliefs Π, $\succeq_A$ is more (less) ambiguity averse than $\succeq_B$ if and only if $\alpha_A \geq (\leq)\alpha_B$, and u_A and u_B are equal up to an affine transformation (GMM, Prop. 12). Furthermore if Π is "centrally symmetric", in that there exists a probability measure π^* – known as the 'centre' of Π – such that $\pi \in \Pi$ if and only if $\pi^* - (\pi - \pi^*) \in \Pi$ (Jewitt and Mukerji, 2012), then the preference $\succeq$ is ambiguity neutral if and only if $\alpha = 0.5$. It follows that, for this centrally symmetric Π, $\succeq$ is ambiguity averse (seeking) if and only if $\alpha > (<)0.5$. In the extreme case where $\alpha = 1$, the preference is equivalent to maxmin expected utility in Gilboa and Schmeidler (1989).

α-MEU is not the only way to characterise choice in a manner that distinguishes the decision-maker's beliefs and preferences towards ambiguity. An alternative is Klibanoff, Marinacci and Mukerji's (2005, KMM) "smooth" representation of choice under ambiguity, and indeed Jewitt and Mukerji (2012) apply their definitions of more ambiguous to both the GMM and KMM representations. We will focus on the α-MEU representation here, since it is very similar in structure to our capital-setting rule (2), which in turn allows for the derivation of particularly tractable expressions for the price of contracts affected by ambiguity, developed in Section 4.

Note that (5) does not constrain the preferences of agents over unbounded books. This means that if we were to define "more ambiguous" in terms of the choices of all ambiguity-averse, -neutral, and -seeking agents with preferences consistent with GMM's representation, we would never be able to describe one unbounded book as being more or less ambiguous than

7

another. And as we wish to do just this, we characterise "more ambiguous" relative to a narrower class of preferences than those consistent with GMM's representation. We thus define $\mathcal{P}_\Pi$ as the set of all preferences over $\mathcal{F}$ that are consistent with GMM's representation, that rank any $f, g \in \mathcal{F}$ according to (5) provided all the expectations in (5) are defined, and that share beliefs given by Π. The ambiguity attitude of any $\succeq \in \mathcal{P}_\Pi$ is determined by α just as in GMM's representation. We say $\succeq \in \mathcal{P}_\Pi$ is f-*constrained* if and only if, given the function u and weight α associated with $\succeq$, $\alpha \cdot \min_{\pi \in \Pi} \int_S u\left(f(s)\right) d\pi + (1 - \alpha) \cdot \max_{\pi \in \Pi} \int_S u\left(f(s)\right) d\pi$ is defined.

Let us now define what it means for one book to be "more ambiguous" than another. We denote the symmetric component of $\succeq$ using $\sim$ as usual.

Definition 1. For any $f, g \in \mathcal{F}$, f is $\mathcal{P}_\Pi$-more ambiguous than g iff:

[i.] For all f- and g-constrained, ambiguity-neutral $\succeq \in \mathcal{P}_\Pi$, $f \sim g$;

[ii.] For any f- and g-constrained $\succeq_A, \succeq_B \in \mathcal{P}_\Pi$, where $\succeq_A$ is ambiguity neutral: if $\succeq_B$ is more ambiguity averse than $\succeq_A$, $g \succeq_B f$; and if $\succeq_A$ is more ambiguity averse than $\succeq_B$, $f \succeq_B g$.

Where the particular configuration of beliefs is unimportant or obvious from the context, we will omit the qualification "$\mathcal{P}_\Pi$-" and simply say f is "more ambiguous" than g. We describe book f as *unambiguous* if, for all $g \in \mathcal{F}$, either g is more ambiguous than f, or f is not more ambiguous than g. f can only be unambiguous if $P_f^\pi = P_f^{\pi'}$ for all $\pi, \pi' \in \Pi$. An ambiguous book is then any book that is not unambiguous.

To state our first results we require some further terminology. First, a Markov kernel from $(\Pi, \mathcal{B}_\Pi)$ to itself is any map $(\pi, E) \mapsto K_\pi(E)$ such that K_π is a probability measure on $\mathcal{B}_\Pi$. For any pair of books f and g, we say K π-*garbles* f into g whenever, for all $E \in \mathcal{B}_\mathbb{R}$, the following condition holds for all $\pi' \in \Pi$:

$$P_g^{\pi'}(E) = \int_\Pi P_f^\pi(E) dK_{\pi'} \qquad (6)$$

The existence of a π-garbling from f to g implies that the likelihood that g pays out in E conditional on any π' is a weighted average of the likelihood that f pays E across all $\pi \in \Pi$. In this sense, f's payoff depends more sensitively on the realisation of the true probability model than g's does.

Let Π be compact, convex and centrally symmetric with centre π^*. Then Jewitt and Mukerji characterise a Markov kernel K from $(\Pi, \mathcal{B}_\Pi)$ to itself as *centre-preserving* if and only if, for all $E \in \mathcal{B}_\mathbb{R}$,

$$P_f^{\pi^*}(E) = \int_\Pi P_f^\pi(E) dK_{\pi^*}$$

Whenever there is a centre-preserving π-garbling from f into g, $P_f^{\pi^*}(E) = P_g^{\pi^*}(E)$ for all $E \in \mathcal{B}_{\mathbb{R}}$. Thus a centre-preserving π-garbling is analogous to a mean-preserving spread familiar from the analysis of risk: just as a mean-preserving spread preserves the expected payoff of a prospect, but makes this payoff more sensitive to the true state of the world, a centre-preserving π-garbling preserves a book's payoff-distribution at the centre, $\pi = \pi^*$, but makes the payoff-distribution more strongly dependent on the true model in Π.

Given this, the first result we report from Jewitt and Mukerji should not come as a surprise.

Proposition 1. *[Jewitt-Mukerji 1] For any $f, g \in \mathcal{F}$, if there is a centre-preserving π-garbling from f to g then f is more ambiguous than g.*

It thus follows that if there is a centre-preserving π-garbling from f into g, then any ambiguity-averse agent whose preferences belong to $\mathcal{P}_{\Pi}$ would prefer f to g. The next step is to show that, given the sufficient condition set out in Proposition 1, a more ambiguous book f incurs a higher (lower) capital holding than its counterpart g if and only if $\hat{\alpha} \geq (\leq)0.5$.

Proposition 2. *Suppose Z_f and Z_g are well defined by (2). Then if there is a centre-preserving π-garbling from f to g, $Z_f \geq (\leq)Z_g$ if and only if $\hat{\alpha} \geq (\leq)0.5$.*

Proof: See Appendix.

Proposition 2 shows that our capital-setting rule encodes ambiguity attitude through the parameter $\hat{\alpha}$ in a parallel manner to the parameter α in the GMM representation. It implies that, all else being equal, an ambiguity-averse insurer, represented by $\hat{\alpha} > 0.5$, will hold a larger amount of capital against the risk of ruin of a more ambiguous book, where the definition of "more ambiguous" comes from Proposition 1. By contrast, an ambiguity-neutral insurer with $\hat{\alpha} = 0.5$ will hold neither more nor less capital, while an ambiguity-seeking insurer ($\hat{\alpha} < 0.5$) will hold less capital.

In turn, we can use this result to consider how ambiguity attitude affects the premium charged for a specific contract:

Corollary 1. *Let $f + c = f'$ and $f + c' = f''$, and suppose there is a centre-preserving π-garbling from f' to f''. Then on the assumption that $L_c = L_{c'}$ and $y > 0$, $p_c > p_{c'}$ for any insurer setting its capital holding according to (2), with $\hat{\alpha} > 0.5$, and its premium price according to (4).*

The Corollary considers the case where the addition of a new contract c to the existing book f results in a more ambiguous book f', compared with the addition of an alternative new contract c', which results in book f''.

The most straightforward reason for this would be that c is in itself more ambiguous than c', although the difference between the ambiguity of the resulting books f' and f'' could also stem from ambiguity about how the payoffs from the new contracts co-vary with the existing book. The assumption that $L_c = L_{c'}$ amounts to a condition that the administrative costs of c and c' be the same (since both c and c' have the same expected loss), while $y > 0$ will surely hold. The Corollary thus tells us that, when setting premiums for any pair of new contracts, an ambiguity-averse insurer will charge a higher premium for the contract that results in a more ambiguous insurance portfolio, all else being equal.

In Section 4 we explore circumstances in which book f' is more ambiguous than f under various distributional assumptions.

3.1 U-Comonotonicity

Proposition 1 applies to any pair of books under any set of beliefs, but it provides only a sufficient condition for one book to be more ambiguous than another. Thus, Proposition 2 does not establish that (2) encodes ambiguity aversion over *all* pairs of books. However, we can use a second result from Jewitt and Mukerji's analysis that provides sufficient and necessary conditions for book f to be more ambiguous than g, provided f, g and Π satisfy a certain condition – known as *U-comonotonicity* – in relation to each other:[5]

Definition 2. Π is U-comonotone for $\mathcal{F}^* \subset \mathcal{F}$ iff Π can be placed in a linear order $\leq_U$ such that for all non-decreasing bounded functions u:

$$\pi \leq_U \pi' \iff \int_S u(f(s))d\pi \leq \int_S u(f(s))d\pi' \text{ for all } f \in \mathcal{F}^*$$

In words, Π is U-comonotone over $\mathcal{F}^*$ if all expected utility maximisers with bounded utility non-decreasing in money and a book belonging to $\mathcal{F}^*$ would agree on a single ranking of which of any pair in Π represented "better news" about the true probability model. This might be the case, for example, where the set $\mathcal{F}^*$ consisted of books that paid out a fixed sum in case of an extreme weather event: Π could then be ordered such that $\pi \leq_U \pi'$ if and only if π' places a lower probability on the extreme weather event than π does. Indeed, this example indicates when it might be plausible to assume U-comonotonicity, namely when each of the set of acts under consideration is stochastically 'similar', in the sense that the realisation of a $\pi \in \Pi$ has similar consequences for the likelihood of the acts in $\mathcal{F}^*$ producing "good" or "adverse" consequences. Jewitt and Mukerji give the example of a pair of bets on the S&P equities index as being stochastically similar (events),

---

[5]Note this is a special case of a more general definition, which can be found in Jewitt and Mukerji (2012).

in comparison with a pair of bets, where one is on the S&P and the other is on the outcome of a boxing match. The acts in this paper are books of insurance and may therefore be considered stochastically similar, especially since the principal interpretation of Π that we offer is of competing estimates of catastrophic risks, which should affect many books' pay-offs in a similar fashion.

To state Jewitt and Mukerji's characterisation of ambiguity aversion under U-comonotonicity, we require a definition of the comparative ambiguity of *events* $E, E' \in \mathcal{B}_{\mathbb{R}}$, in addition to the definition of the comparative ambiguity of acts given previously in Proposition 1. For any pair of payoffs x and y, xEy denotes the binary act that pays x if the realised state $s \in E$ and y otherwise.

Definition 3. Given events $E, E' \in \mathcal{B}_{\mathbb{R}}$, E is a more ambiguous event than E' if, for all ambiguity neutral $\succeq_A \in \mathcal{P}$,

$$xE'y \sim_A xEy \text{ and } x(\neg E')y \sim_A x(\neg E)y;$$

for all $\succeq_B \in \mathcal{P}$, such that $\succeq_B$ is more ambiguity averse than $\succeq_A$,

$$xE'y \succeq_B xEy \text{ and } x(\neg E')y \succeq_B x(\neg E)y;$$

for all $\succeq_B \in \mathcal{P}$, such that $\succeq_A$ is more ambiguity averse than $\succeq_B$,

$$xE'y \preceq_B xEy \text{ and } x(\neg E')y \preceq_B x(\neg E)y,$$

where $x > y$.

In the specific context of α-MEU preferences, E is a more ambiguous event than E' if and only if E' is a centre-preserving π-garbling of E for Π compact, convex and centrally symmetric, i.e. where acts f and g in (6) are unit bets on E and E' respectively.

The following Proposition characterises ambiguity aversion under U-comonotonicity, by establishing that events constituting adverse payoffs under book f are more ambiguous events than the corresponding adverse payoffs under book g.[6]

Proposition 3. *[Jewitt-Mukerji 2] Suppose Π is compact, convex and centrally symmetric with centre $\pi*$ and is U-comonotone on $\{f, g\}$. Then the following three statements are equivalent:*

---

[6]The result reported here is slightly different to that in Jewitt and Mukerji, who define U-comonotonicity in terms all non-decreasing (bounded and unbounded) utility functions but consider only bounded books. Our statement of the result encompasses all books but defines U-comonotonicity in terms of bounded utility functions; the proof is nonetheless as in Jewitt and Mukerji with obvious modifications.

[1.] f is $\mathcal{P}_\Pi$-more ambiguous than g;

[2.] For each x, $E_f^x \equiv \{s \in S : f(s) \leq x\}$, $E_g^x \equiv \{s \in S : g(s) \leq x\} \in \mathcal{B}_\mathbb{R}$, E_f^x is a more ambiguous event than E_g^x;

[3.] There is a centre-preserving π-garbling from f to g, for $\pi, \pi' \in \Pi$, $\pi \leq_U \pi'$, the map $(\alpha, h) \mapsto P_h^{\alpha\pi+(1-\alpha)\pi'}$ is supermodular on $[0,1] \times \{f,g\}$. Specifically for $0 \leq \alpha < \alpha' \leq 1$,

$$P_g^{\alpha\pi+(1-\alpha)\pi'} - P_g^{\alpha'\pi+(1-\alpha')\pi'} \leq P_f^{\alpha\pi+(1-\alpha)\pi'} - P_f^{\alpha'\pi+(1-\alpha')\pi'}$$

The intuition behind Proposition 3 is that if f is more ambiguous than g, its payoff distribution is more sensitive to the realisation of the true model in Π. This result allows us to obtain the equivalent of Proposition 2.

Proposition 4. *If Π is U-comonotone on $\{f, g\}$ and f is more ambiguous than g then $Z_f \geq (\leq)Z_g$ if and only if $\hat{\alpha} \geq (\leq)0.5$.*

Proof: see Appendix.

Proposition 4 has the same implications for the capital holdings of insurers as Proposition 2, the difference being in the way in which any book of insurance is defined as being more ambiguous than another. An ambiguity-averse insurer with $\hat{\alpha} > 0.5$ will hold a larger amount of capital against the risk of ruin of the more ambiguous book in the pair, all else being equal. In addition, Proposition 4 has the same implications for the premium price attached to a new contract that results in a more ambiguous insurance portfolio. It will be higher than the premium charged for a new contract that results in a less ambiguous insurance portfolio, defined according to Proposition 3, if and only if the insurer is ambiguity-averse with $\hat{\alpha} > 0.5$.

4 Contract pricing under ambiguity

We now examine the impact of ambiguity on the price of an individual contract given a capital holding set according to (2) and competitive pricing of premiums according to (4). We derive several pricing formulae, which show explicitly how introducing ambiguity leads to a departure from a benchmark pricing formula in the absence of ambiguity, i.e. we explicitly identify an additional 'ambiguity load'. Our starting point is a model, influential in the actuarial literature and in the insurance industry, which utilises information about the mean and variance of losses on the new contract, as well as on the existing book (Kreps, 1990).

We examine four types of ambiguity, chosen on the basis of their analytical tractability and applicability to real insurance problems. The four cases differ on the distribution of model parameters assumed under the measure ν.

We only consider contract pricing, thus we ignore deductibles, co-insurance and other design options that an insurer might use to manage ambiguity. These would be interesting avenues for future research (see Amarante et al., 2015).

4.1 Benchmark: no ambiguity

As a benchmark for what follows we review the case where the insurer's information is unambiguous. The set of books under consideration is $\mathcal{F}_0 \subset \mathcal{F}$, where $\mathcal{F}_0$ is defined relative to a given Π as follows: $f \in \mathcal{F}_0$ iff the density of $f(s)$ under P_f^π on $\{S, \mathcal{B}\}$ for all $\pi \in \Pi$ is parameterised by mean μ_f and variance σ_f^2. We define the addition operation over $\mathcal{F}_0$ pointwise – that is, for $f, f' \in \mathcal{F}_0$, $f + f' = f''$ where $f''(s) = f(s) + f'(s)$ for all s – and note that $\mathcal{F}_0$ is closed under addition – i.e. if $f, f' \in \mathcal{F}_0$, $f + f' \in \mathcal{F}_0$.

It is worth emphasising that even in this framework we assume that the insurer sets its capital holding according to rule (2) and we allow Π to be non-singleton – implying that, across the class of all books, the insurer may face some ambiguity. However, because we restrict our focus in the benchmark case to $\mathcal{F}_0$, the insurer faces no ambiguity and therefore its capital holding rule is equivalent to that in (1). We adopt this approach in order to make clearer the generalisations that follow to richer sets of books.

Where Φ is a standardised cdf, given whatever assumption about functional form the insurer finds appropriate (e.g. normal, gamma, etc.), and $-z = \Phi^{-1}(\theta)$, the insurer's capital holding for $f \in \mathcal{F}_0$ is determined by:

$$Z_f = z\sigma_f - \mu_f \tag{7}$$

Given competitive pricing (3) this implies:

$$\mu_c = \frac{yz}{(1+y)}(\sigma_{f'} - \sigma_f)$$

Recalling that where $\rho_{c,f}$ is the correlation coefficient for the random variables $c(s)$ and $f(s)$, $\sigma_{f'}^2 = \sigma_f^2 + \sigma_c^2 + 2\sigma_c\sigma_f\rho_{c,f}$, we have:

$$\sigma_{f'} - \sigma_f = \sigma_c \frac{2\sigma_f\rho_{c,f} + \sigma_c}{\sigma_{f'} + \sigma_f}$$

And hence, where $\mathcal{R}_{c,f} := (yz/(1+y))(2\sigma_f\rho_{c,f} + \sigma_c)/(\sigma_{f'} + \sigma_f)$:

$$\mu_c = \mathcal{R}_{c,f}\sigma_c$$

Using the general expression for the price of a premium set by a competitive insurer, we can now state Kreps's (1990) more specific pricing result, the proof of which is immediate from the analysis above:

Proposition 5. *[Pricing without ambiguity] If $f, c \in \mathcal{F}_0$, then a competitive insurer with book f will set p_c as follows:*

$$p_c = L_c + \mathcal{R}_{c,f} \sigma_c \qquad (8)$$

The second element on the right-hand side is the *risk load* for contract c. Note that it arises solely as a consequence of the insurer's need to limit the probability of ruin to a certain level (encoded in rule (2)): without this constraint the competitive price of the contract would simply be L_c. As one would expect, the risk load is increasing in the riskiness of the contract (measured by σ_c), the contract's correlation with the insurer's pre-existing book ($\rho_{c,f}$), the opportunity cost of capital (y), and it is decreasing in the acceptable probability of loss (increasing in z – a decreasing function of θ).

4.2 Mean uncertain; variance known

4.2.1 Mean uniformly distributed

Our first case of ambiguity involves considering a space of books $\mathcal{F}_1$ that, given some Π and ν, satisfies[7]: (1.i) for all $f \in \mathcal{F}_1$ and all $\pi \in \Pi$, $f(s)$ under P^π on $\{S, \mathcal{B}\}$ has mean μ_f^π and variance σ_f^2 ; (1.ii) for all $f \in \mathcal{F}_1$, μ_f^π is uniformly distributed on $[a_f, b_f]$ given ν on $\{\Pi, \mathcal{B}_\Pi\}$; (1.iii) $\mathcal{F}_1$ is closed under addition; and (1.iv) $\mathcal{F}_0 \subseteq \mathcal{F}_1$. Note that it is impossible to satisfy the additivity condition without violating (1.ii) unless, for all $f, f' \in \mathcal{F}_1$:

$$\mu_{f'}^\pi = a_{f'} + (\mu_f^\pi - a_f) \frac{(b_{f'} - a_{f'})}{(b_f - a_f)} \qquad (9)$$

which implies that Π is U-comonotone for $\mathcal{F}_1$. In the cases examined here, a more ambiguous book therefore incurs a higher capital holding as per Propositions 3 and 4 in Section 3.

To illustrate where a structure like this might apply, consider the following example.

Example 1. Suppose our insurer has a collection of forecasts at its disposal, all of which agree on the payoff-variance of any given book, but amongst which there is disagreement over certain books' payoff-expectations. Specifically, there is a most pessimistic simulation, which reports the lowest mean payoff for all the books – for book f this is a_f – and a most optimistic simulation, which gives the highest reported mean for any book – b_f for book f. For any book, it is sure that the variance is as reported – σ_f^2 for book f – and thinks the true mean must lie somewhere between these optimistic and

---

[7]Note that $\mathcal{F}_1$ may not be unique given Π and ν. This is also the case for $\mathcal{F}_2, \mathcal{F}_3$, and $\mathcal{F}_4$ introduced below.

pessimistic bounds. It constructs Π and ν using two assumptions. First, the members of Π are ordered according to their pessimism so that (9) is satisfied and for any $c \in [a_f, b_f]$, $\mu_f^\pi = c$ for one $\pi \in \Pi$. Second, ν is set such that condition (1.ii), imposing a uniform distribution on μ_f^π, holds. The first assumption may be justified in case the insurer finds it reasonable, while the second is reasonable provided it has no evidence to suggest any value of μ_f^π in $[a_f, b_f]$ is more plausible than any other, in which case the uniformity of μ_f^π follows from the principle of insufficient reason. Under these assumptions, any book it considers belongs to $\mathcal{F}_1$ given Π and ν.

Given the decision rule (2), Z_f is set to reduce the probability of ruin to an acceptable level under the weighted sum of the most pessimistic and optimistic models, i.e. for any $f \in \mathcal{F}_1$

$$Z_f = z\sigma_f - [\hat{\alpha} \cdot a_f + (1 - \hat{\alpha}) \cdot b_f]$$
$$= z\sigma_f - \hat{\alpha}a_f + (\hat{\alpha} - 1)b_f$$

Where the models are uniformly distributed this implies

$$Z_f = z\sigma_f + (2\hat{\alpha} - 1) \cdot \left(\sqrt{3}\text{sd}[\mu_f^\pi]\right) - \mu_f \qquad (10)$$

where $\text{sd}[\mu_f^\pi]$ is the standard deviation of the random variable μ_f^π, equal to $\sqrt{1/12(b_f - a_f)^2}$ under the uniformity assumption.

We now proceed in parallel to the exposition of the previous sub-section, supposing that a competitive insurer with book $f \in \mathcal{F}_1$ accepts the further contract $c \in \mathcal{F}_1$ and thereby ends up with the book $f' = f + c$. Using (10) and (3) as above, we obtain:

$$\mu_c = \mathcal{R}_{c,f}\sigma_c + \frac{\sqrt{3}y(2\hat{\alpha} - 1)}{1 + y}\left(\text{sd}[\mu_{f'}^\pi] - \text{sd}[\mu_f^\pi]\right)$$

Using the fact that

$$\text{Var}[\mu_{f'}^\pi] = \text{Var}[\mu_f{}^\pi] + \text{Var}[\mu_c^\pi] + 2\text{Cov}[\mu_f^\pi, \mu_c^\pi]$$

and furthermore that, given (9), $\text{sd}[\mu_f^\pi]\text{sd}[\mu_c^\pi] = \text{Cov}\left[\mu_f^\pi, \mu_c^\pi\right]$, we can define

$$\mathcal{A}_{c,f,1} \equiv \left[\frac{\sqrt{3}y(2\hat{\alpha} - 1)}{1 + y}\right]\left(\frac{2\text{sd}[\mu_f^\pi] + \text{sd}[\mu_c^\pi]}{\text{sd}[\mu_{f'}^\pi] + \text{sd}[\mu_f^\pi]}\right)$$

so that our first pricing result under ambiguity follows straightforwardly:

Proposition 6. *[Pricing with uniform mean] If $f, c \in \mathcal{F}_1$, then a competitive insurer with book f will set*

$$p_c \quad = \quad L_c + \mathcal{R}_{c,f}\sigma_c + \mathcal{A}_{c,f,1}\text{sd}[\mu_c^\pi]$$

It is easy to see how Proposition 6 generalises Proposition 5. If c is unambiguous then it must belong to $\mathcal{F}_0$, in which case $\mathrm{sd}[\mu_c^\pi] = 0$ and so p_c is set according to (8). However, if c is in $\mathcal{F}_1 \setminus \mathcal{F}_0$ – that is to say c has ambiguous returns – then p_c also incorporates an *ambiguity load* equal to $\mathcal{A}_{c,f,1}\mathrm{sd}[\mu_c^\pi]$. The ambiguity load is positive, provided the index of ambiguity aversion $\hat{\alpha} > 0.5$. It is increasing in $\hat{\alpha}$ and in $\mathrm{sd}[\mu_c^\pi]$, the (approximate) measure of ambiguity in c[8]. It is also increasing in the cost of capital, and it is increasing in the ambiguity of the pre-existing book (measured by $\mathrm{sd}[\mu_f^\pi]$) whenever $\hat{\alpha} > 0.5$.

4.2.2 Mean triangularly distributed

We now consider an alternative space of books, $\mathcal{F}_2$, defined such that given Π and ν: (2.i) for all $f \in \mathcal{F}_2$ and all $\pi \in \Pi$, $f(s)$ under P_f^π on $\{S, \mathcal{B}\}$ has mean μ_f^π and variance σ_f^2 ; (2.ii) for all $f \in \mathcal{F}_2$, μ_f^π has a symmetric triangular distribution on $[a_f, b_f]$ given ν on $\{\Pi, \mathcal{B}_\Pi\}$; (2.iii) $\mathcal{F}_2$ is approximately closed under addition[9]; and (2.iv) $\mathcal{F}_0 \subseteq \mathcal{F}_2$. Conditions (2.i), (2.iii), and (2.iv) mirror their counterparts in the analysis of a uniform mean. Once again, (2.ii) and (2.iii) may only be satisfied when (9) holds for all $f, f' \in \mathcal{F}_2$ and Π is U-comonotone for $\mathcal{F}_2$.

To illustrate the applicability of $\mathcal{F}_2$, we extend Example 1.

Example 2. Suppose the insurer from Example 1 thinks that, for any book f, values of μ_f^π closer to the midpoint of the range $[a_f, b_f]$ are more probable than those further away from it, i.e. roughly speaking that models with more extreme forecasts of the mean loss are less likely to be correct. Provided these beliefs are reasonably approximated by the assumption that μ_f^π is triangularly distributed[10] with minimum a_f, maximum b_f and mode $(a_f + b_f)/2$, it might proceed by again assuming the members of Π are ordered according to their pessimism and by setting ν so that (2.ii) is satisfied. Given Π and ν thus constructed, every book it considers will belong to $\mathcal{F}_2$.

For any $f \in \mathcal{F}_2$ we have:

$$Z_f \;\; = \;\; z\sigma_f + (2\hat{\alpha} - 1) \cdot \left(\sqrt{6}\mathrm{sd}[\mu_f^\pi] \right) - \mu_f$$

---

[8]See discussions on this point in Jewitt and Mukerji (2012) and Maccheroni et al. (2010).

[9]Triangular distributions are not closed under addition, so what this requires is that an insurer finds it appropriate to approximate the sum of two triangular distributions $f, f' \in \mathcal{F}_2$ with another triangular distribution that is a member of $\mathcal{F}_2$.

[10]We choose this distribution as (2) does not well-define the capital holding unless μ_f^π has a bounded support, however triangular distributions are also frequently used to characterise subjective probability distributions in probability-elicitation

which, for $f' = f + c$ and $f, c \in \mathcal{F}_2$, yields:

$$\mu_c = \mathcal{R}_{c,f}\sigma_c + \frac{\sqrt{6}y(2\hat{\alpha} - 1)}{1 + y}\left(\text{sd}[\mu_{f'}^\pi] - \text{sd}[\mu_f^\pi]\right)$$

Where we use $\mathcal{A}_{c,f,2}$ to denote $(\frac{\sqrt{6}y(2\hat{\alpha}-1)}{1+y})\left(\frac{2\text{sd}[\mu_f^\pi]+\text{sd}[\mu_c^\pi]}{\text{sd}[\mu_{f'}^\pi]+\text{sd}[\mu_f^\pi]}\right)$, this gives us:

$$\mu_c = \mathcal{R}_{c,f}\sigma_c + \mathcal{A}_{c,f,2}\text{sd}[\mu_c^\pi]$$

from which the next pricing result is immediate:

Proposition 7. *[Pricing with triangular mean] If $f, c \in \mathcal{F}_2$, then a competitive insurer with book f will set*

$$p_c = L_c + \mathcal{R}_{c,f}\sigma_c + \mathcal{A}_{c,f,2}sd[\mu_c^\pi]$$

The result generalises Proposition 5 by incorporating an ambiguity load that is zero for $c \in \mathcal{F}_0$ and increasing in $\text{sd}[\mu_c^\pi]$. Like Proposition 6 it is also increasing in the index of ambiguity aversion $\hat{\alpha}$ and in the cost of capital, while the relationship between the ambiguity load and the ambiguity of the existing book is the same as before. However, the switch from a uniform distribution to a triangular distribution implies that, for any given standard deviation, the range of possible values increases, thus an ambiguity-averse reinsurer would require a greater ambiguity load.

4.3   Mean known; variance uncertain

We now focus on a space of books, $\mathcal{F}_3$, defined for a given Π and ν such that: (3.i) for all $f \in \mathcal{F}_3$ and all $\pi \in \Pi$, $f(s)$ under P_f^π on $\{S, \mathcal{B}\}$ has mean μ_f and variance $(\sigma_f^\pi)^2$; (3.ii) for all $f \in \mathcal{F}_3$, σ_f^π has a uniform distribution on $[a_f, b_f]$ given ν on $\{\Pi, \mathcal{B}_\Pi\}$; (3.iii) $\mathcal{F}_3$ is closed under addition; and (3.iv) $\mathcal{F}_0 \subseteq \mathcal{F}_3$. As in previous sections, additivity and the uniformity of σ_f^π imply that for any $f, f' \in \mathcal{F}_3$ and $\pi \in \Pi$, σ_f^π and $\sigma_{f'}^\pi$ are linearly related as follows:

$$\sigma_{f'}^\pi = a_{f'} + (\sigma_f^\pi - a_f)\frac{(b_{f'} - a_{f'})}{(a_f - b_f)} \tag{11}$$

Unless $\mathcal{F}_3 = \mathcal{F}_0$, Π is not U-comonotone for $\mathcal{F}_3$.

We imagine this case applying to an insurer in an analogous position to that described by Example 1, except with a range of estimates of the standard deviation of losses and certainty over the mean.[11]

---

[11]Though note the appeal to the principle of insufficient reason to justify the uniformity of σ_f^π for all f is weaker here. The insurer could equally invoke the principle to impose the uniformity of $(\sigma_f^\pi)^2$, in which case the collection of books it considers could not satisfy (3.ii).

Working as before, Z_f is set to reduce the probability of ruin to an acceptable level under the weighted sum of the most pessimistic and optimistic models, i.e. for any $f \in \mathcal{F}_3$

$$
\begin{aligned}
Z_f &= \hat{\alpha} \cdot z \cdot b_f + (1 - \hat{\alpha}) \cdot z \cdot a_f - \mu_f \\
&= z \left[\hat{\alpha} b_f + (1 - \hat{\alpha}) a_f \right] - \mu_f \\
&= z \left[\mathbb{E}[\sigma_f^\pi] + (2\hat{\alpha} - 1) \cdot \left(\sqrt{3} \mathrm{sd}[\sigma_f^\pi] \right) \right] - \mu_f
\end{aligned} \tag{12}
$$

and thus, where $f, c \in \mathcal{F}_3$ and $f' = f + c$, (3) implies:

$$
\mu_c = \frac{yz}{1 + y} \left(E[\sigma_{f'}^\pi - \sigma_f^\pi] \right) + \frac{\sqrt{3} yz (2\hat{\alpha} - 1)}{1 + y} \left(\mathrm{sd}[\sigma_{f'}^\pi] - \mathrm{sd}[\sigma_f^\pi] \right)
$$

Using the fact that $\mathbb{E}[\sigma_{f'}^\pi - \sigma_f^\pi] = \mathbb{E}[\sigma_c^\pi \frac{2\sigma_f^\pi \rho_{cf} + \sigma_c^\pi}{\sigma_{f'}^\pi + \sigma_f^\pi}]$, we can further decompose the risk load $\frac{yz}{1+y} \left(E[\sigma_{f'}^\pi - \sigma_f^\pi] \right)$ into two terms, first recovering the equivalent of the risk load in the absence of ambiguity, and second obtaining a term capturing how the risk load depends on ambiguity over σ_c^π:

$$
\begin{aligned}
\mu_c &= \mathbb{E}[\mathcal{R}_{c,f} \sigma_c^\pi] + \frac{\sqrt{3} yz (2\hat{\alpha} - 1)}{1 + y} \left(\mathrm{sd}[\sigma_{f'}^\pi] - \mathrm{sd}[\sigma_f^\pi] \right) \\
&= \mathbb{E}[\sigma_c^\pi] \mathbb{E}[\mathcal{R}_{c,f}] + \mathrm{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}] + \frac{\sqrt{3} yz (2\hat{\alpha} - 1)}{1 + y} \left(\mathrm{sd}[\sigma_{f'}^\pi] - \mathrm{sd}[\sigma_f^\pi] \right)
\end{aligned}
$$

To obtain the ambiguity load, take a similar approach as before, using the fact that $\mathrm{Var}[\sigma_{f'}^\pi] = \mathrm{Var}[\sigma_f^\pi] + \mathrm{Var}[\sigma_c^\pi] + 2\mathrm{sd}[\sigma_f^\pi]\mathrm{sd}[\sigma_c^\pi]\mathrm{corr}\left[\sigma_f^\pi, \sigma_c^\pi\right]$ and, given (11), that $\mathrm{sd}[\sigma_f^\pi]\mathrm{sd}[\sigma_c^\pi] = \mathrm{Cov}\left[\sigma_f^\pi, \sigma_c^\pi\right]$. Thus, defining

$$
\mathcal{A}_{c,f,3} = \frac{\sqrt{3} yz (2\hat{\alpha} - 1)}{1 + y} \left(\frac{2\mathrm{sd}[\sigma_f^\pi] + \mathrm{sd}[\sigma_c^\pi]}{\mathrm{sd}[\sigma_f^\pi] + \mathrm{sd}[\sigma_{f'}^\pi]} \right)
$$

we have

$$
\mu_c = \mathbb{E}[\sigma_c^\pi] \mathbb{E}[\mathcal{R}_{c,f}] + \mathrm{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}] + \mathcal{A}_{c,f,3} \mathrm{sd}[\sigma_c^\pi]
$$

which gives us our next pricing result:

Proposition 8. *[Pricing with uniform standard deviation] If $f, c \in \mathcal{F}_3$, then a competitive insurer with book f will set*

$$
p_c = L_c + \mathbb{E}[\sigma_c^\pi] \mathbb{E}[\mathcal{R}_{c,f}] + \mathrm{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}] + \mathcal{A}_{c,f,3} \mathrm{sd}[\sigma_c^\pi]
$$

Once again, whenever $c \in \mathcal{F}_0$, the pricing formula above reduces to (8). In contrast to our previous results, however, introducing ambiguity affects the

price of a contract via two additional terms rather than one. First, as in our earlier results, there is a term, $\mathcal{A}_{c,f,3}\text{sd}[\sigma_c^\pi]$, that is increasing in the ambiguity of c, $\text{sd}[\sigma_c^\pi]$, in the index of ambiguity aversion $\hat{\alpha}$, in the cost of capital and the smaller is the acceptable probability of loss (the larger is z). The dependence of the ambiguity load on z is new and follows immediately from (12) – it is due to the fact that ambiguity in this example concerns the variance of returns, rather than mean returns. The second term, $\text{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}]$, reflects the fact that uncertainty over σ_c^π leads to uncertainty over the risk load. Since $\text{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}]$ could in principle depend negatively on the ambiguity of c, the overall ambiguity load could, in contrast to the other cases examined so far, be negative.

4.4 Mean and variance uncertain

As a final exercise, we consider an informational structure that nests two of the cases described above: where both the mean and variance are independently uniformly distributed. Thus we consider a space of books, $\mathcal{F}_4$, that satisfies: (4.i) for all $f \in \mathcal{F}_4$ and all $\pi \in \Pi$, $f(s)$ under P_f^π on $\{S, \mathcal{B}\}$ has mean μ_f^π and variance $\left(\sigma_f^\pi\right)^2$; (4.ii) for all $f \in \mathcal{F}_4$, μ_f^π is uniformly distributed on $[a_f, b_f]$, σ_f^π is uniformly distributed on $[a'_f, b'_f]$, and μ_f^π and σ_f^π are independent given ν on $\{\Pi, \mathcal{B}_\Pi\}$, ; (4.iii) $\mathcal{F}_4$ is closed under addition; and (4.iv) $\mathcal{F}_0 \subseteq \mathcal{F}_4$. Given this definition, for any pair $f, f' \in \mathcal{F}_4$, μ_f^π and σ_f^π must satisfy conditions (9) and (11) (the latter with obvious relabelling). Apart from cases where $\mathcal{F}_4 \in \{\mathcal{F}_0, \mathcal{F}_1\}$, Π is not U-comonotone for $\mathcal{F}_4$.

Proceeding in the usual way, we have, for any $f \in \mathcal{F}_4$:

$$Z_f = z\left[\mathbb{E}[\sigma_f^\pi] + (2\hat{\alpha} - 1) \cdot \left(\sqrt{3}\text{sd}[\sigma_f^\pi]\right)\right] + (2\hat{\alpha} - 1) \cdot \left(\sqrt{3}\text{sd}[\mu_f^\pi]\right) - \mu_f$$

So for $f, c \in \mathcal{F}_4$, a competitive insurer with book f prices c such that

$$\mu_c = \frac{yz}{1+y}\left(E[\sigma_{f'}^\pi - \sigma_f^\pi]\right) + \frac{\sqrt{3}yz(2\hat{\alpha} - 1)}{1 + y}\left(\text{sd}[\sigma_{f'}^\pi] - \text{sd}[\sigma_f^\pi]\right)$$
$$+ \frac{\sqrt{3}y(2\hat{\alpha}-1)}{1+y}\left(\text{sd}[\mu_{f'}^\pi] - \text{sd}[\mu_f^\pi]\right)$$

It is then clear that we can progress using steps from our analyses of $\mathcal{F}_1$ and $\mathcal{F}_3$ above to reach our final pricing formula.

Proposition 9. *[Pricing with independent uniform mean and standard deviation] Where $f, c \in \mathcal{F}_4$, a competitive insurer with book f will offer*

$$p_c = L_c + \mathbb{E}[\sigma_c^\pi]\mathbb{E}[\mathcal{R}_{c,f}] + \text{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}] + \mathcal{A}_{c,f,1}\text{sd}[\mu_c^\pi] + \mathcal{A}_{c,f,3}\text{sd}[\sigma_c^\pi]$$

Thus, where books and contracts belong to $\mathcal{F}_4$, the ambiguity load for any contract is the sum of a component ($\mathcal{A}_{c,f,1}\text{sd}[\mu_c^\pi]$) arising due to ambiguity in the contract's mean, a component ($\mathcal{A}_{c,f,3}\text{sd}[\sigma_c^\pi]$) reflecting ambiguity

in its standard deviation, and there is also the effect of ambiguity about the standard deviation of the contract on the risk load. That $\mathcal{A}_{c,f,1}\mathrm{sd}[\mu_c^\pi]$ and $\mathcal{A}_{c,f,3}\mathrm{sd}[\sigma_c^\pi]$ are additive results from our restriction that the mean and standard deviation are independent; another way of arriving at the same formula would be to assume that the mean and standard deviation were linearly related, with higher variances corresponding to lower means.

5   Concluding Remarks

The main contribution of this paper has been to establish a clear connection between ambiguity and the pricing of (re)insurance. We show, at a general level, that under our capital-setting rule increasing ambiguity leads to higher capital holdings and thus to higher costs, provided the (re)insurer is averse to ambiguity about the risk of ruin. We then show how, under a range of distributional assumptions, our capital-setting rule gives rise to particular pricing formulae for insurance contracts, all composed of distinct risk and ambiguity loads. These pricing formulae are testable predictions of the theory.

Admittedly we have had to make relatively specific assumptions about the probability distributions describing ambiguity. Since they must be bounded in order that our capital-setting rule is well defined, we employ uniform and triangular distributions. However, these two distributional forms have quite strong appeal as characterisations of ambiguous beliefs. The uniform distribution follows from the application of the principle of insufficient reason, which might often be deemed appropriate, if for some reason (e.g. insufficient data or dependence of different models) the comparative performance of different forecasting models cannot be evaluated. The triangular distribution is also frequently used to characterise subjective probability distributions in probability-elicitation exercises.

But how tenable is our assumption that the capital-setting rule takes the form specified in (2)? From a descriptive perspective, we have already shown that its implications for pricing decisions are consistent with the behavioural evidence in the literature. Of the survey-based studies mentioned in the Introduction, Hogarth and Kunreuther (1992) is distinctive in that it provides tentative evidence from a sample of actuaries of the decision procedures they actually followed. There was some evidence of the use of heuristics to load the premium, such as a simple, *ad hoc* multiplying coefficient on the expected value of the premium, or on the variance of the loss distribution. This is on the face of it at odds with the mechanics of the decision process posited here. At the same time, however, there was also evidence that actuaries had in mind the effect the new contract would have on the overall risk of the insurer's ruin, as in our framework. Indeed, the risk of ruin is known to be an

important consideration more generally when insurers set capital holdings and price contracts, especially for catastrophe risks (e.g. Kunreuther and Michel-Kerjan, 2009). Considering the rule from a normative perspective, one could evaluate the axiomatic foundations of the similar α-MEU rule as set out in Ghirardato et al. (2004). However, decision rules like this are typically motivated from the perspective of an individual decision-maker or social planner – whether it is rational for a corporate entity to follow them remains an open question.

The results in Section 4 suggest insurance contract prices should be increasing in the insurer's degree of ambiguity aversion, and in their ambiguity (as measured by the variance of their uncertain distributional parameters), provided the insurer is ambiguity averse overall. In practice this may not hold if our assumption that models are ordered by their pessimism over the uncertain parameters is violated, for in these cases increasing ambiguity in a contract may allow the insurer to "hedge" against the ambiguity in its pre-existing book. We do not explore this kind of information structure for reasons of tractability and note that, in any case, our assumption is reasonable for some classes of insurance book. For instance, models of the losses arising from natural disasters or terrorism may be ranked according to their pessimism over the likelihood of these events.

References

ALARY, D., C. GOLLIER, AND N. TREICH (2013): "The effect of ambiguity aversion on insurance and self-protection," *Economic Journal*, 123, 1188–1202.

AMARANTE, M., M. GHOSSOUB, AND E. PHELPS (2015): "Ambiguity on the insurer's side: the demand for insurance," *Journal of Mathematical Economics*, 58, 61–78.

BARLOW, C., R. P. NYE, M. G. WHITE, AND H. L. WINCOTE (1993): "Prudential Margins," in *1993 General Insurance Covention*.

BENDER, M., T. KNUTSON, R. TULEYA, J. SIRUTIS, G. VECCHI, S. GARNER, AND I. HELD (2010): "Modeled impact of anthropogenic warming on the frequency of intense Atlantic hurricanes," *Science*, 327, 454.

CABANTOUS, L. (2007): "Ambiguity aversion in the field of insurance: insurers' attitude to imprecise and conflicting probability estimates," *Theory and Decision*, 62, 219–240.

CABANTOUS, L., D. HILTON, H. KUNREUTHER, AND E. MICHEL-KERJAN (2011): "Is imprecise knowledge better than conflicting expertise? Evi-

dence from insurers' decisions in the United States," *Journal of Risk and Uncertainty*, 42, 211–232.

CAIRNS, A. J. (2000): "A discussion of parameter and model uncertainty in insurance," *Insurance: Mathematics and Economics*, 27, 313–330.

ELLSBERG, D. (1961): "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics*, 75, 643–669.

EPSTEIN, L. G. (1999): "A definition of uncertainty aversion," *The Review of Economic Studies*, 66, 579–608.

GARLAPPI, L., R. UPPAL, AND T. WANG (2007): "Portfolio selection with parameter and model uncertainty: a multi-prior approach," *Review of Financial Studies*, 20, 41–81.

GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): "Differentiating ambiguity and ambiguity attitude," *Journal of Economic Theory*, 118, 133–173.

GILBOA, I. AND D. SCHMEIDLER (1989): "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, 18, 141–153.

GOLLIER, C. (2014): "Optimal insurance design of ambiguous risks," *Economic Theory*, 57, 555–576.

HANSEN, L. AND T. SARGENT (2001): "Robust control and model uncertainty," *American Economic Review*, 91, 60–66.

——— (2008): *Robustness*, Princeton University Press.

HOGARTH, R. AND H. KUNREUTHER (1989): "Risk, ambiguity, and insurance," *Journal of Risk and Uncertainty*, 2, 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, 35–60.

JEWITT, I. AND S. MUKERJI (2012): "Ordering ambiguous acts," Discussion paper, Department of Economics, University of Oxford.

KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73, 1849–1892.

KNUTSON, T., J. SIRUTIS, S. GARNER, G. VECCHI, AND I. HELD (2008): "Simulated reduction in Atlantic hurricane frequency under twenty-first-century warming conditions," *Nature Geoscience*, 1, 359–364.

KREPS, R. (1990): "Reinsurer risk loads from marginal surplus requirements," *Proceedings of the Casualty Actuarial Society*, LXXVII, 196–203.

KUNREUTHER, H., R. HOGARTH, AND J. MESZAROS (1993): "Insurer ambiguity and market failure," *Journal of Risk and Uncertainty*, 7, 71–87.

KUNREUTHER, H., J. MESZAROS, R. HOGARTH, AND M. SPRANCA (1995): "Ambiguity and underwriter decision processes," *Journal of Economic Behavior & Organization*, 26, 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.

MACCHERONI, F., M. MARINACCI, AND D. RUFFINO (2010): "Alpha as ambiguity: robust mean-variance portfolio analysis," *IGIER Working Paper 373, Bocconi University.*

MACHINA, M. AND D. SCHMEIDLER (1992): "A more robust definition of subjective probability," *Econometrica*, 60, 745–80.

RANGER, N. AND F. NIEHOERSTER (2012): "Deep uncertainty in long-term hurricane risk: scenario generation and implications for future climate experiments," *Global Environmental Change*, 22, 703–712.

SAVAGE, L. (1954): *The Foundations of Statistics*, John Wiley and Sons.

STONE, J. (1973): "A theory of capacity and the insurance of catastrophe risks (part i)," *Journal of Risk and Insurance*, 231–243.

ZHU, W. (2011): "Ambiguity aversion and an intertemporal equilibrium model of catastrophe-linked securities pricing," *Insurance: Mathematics and Economics*, 49, 38–46.

Appendix

Proof of Proposition 2

Let $\overline{\pi}^g = \arg\max_{\pi \in \Pi} P_g^\pi(-Z_g)$, $\underline{\pi}^g = \arg\min_{\pi \in \Pi} P_g^\pi(-Z_g)$, $\overline{\pi}^f = \arg\max_{\pi \in \Pi} P_g^\pi(-Z_f)$, and $\underline{\pi}^f = \arg\min_{\pi \in \Pi} P_g^\pi(-Z_f)$, and use π^* to denote the centre of Π. Various steps in the following use the fact that, given Π is compact and convex, P_f^π is *mixture linear* in $\pi \in \Pi$ for any $f \in \mathcal{F}$, meaning $P_f^{\lambda\pi+(1-\lambda)\pi'} = \lambda \cdot P_f^\pi + (1-\lambda) \cdot P_f^{\pi'}$ for $\pi, \pi' \in \Pi$, $\lambda \in [0,1]$.

By the central symmetry of Π, there exists a $\pi \in \Pi$ such that

$$P_g^\pi(-Z_g) = P_g^{\pi^*}(-Z_g) - (P_g^{\overline{\pi}^g}(Z_g) - P_g^{\pi^*}(-Z_g)) \tag{13}$$

and there can be no $\pi' \in \Pi$ such that $P_g^{\pi'}(-Z_g) < P_g^\pi(-Z_g)$, since by central symmetry this would entail the existence of $\pi'' \in \Pi$ such that $P_g^{\pi''}(-Z_g) = P_g^{\pi^*}(-Z_g) - (P_g^{\pi'}(-Z_g) - P_g^{\pi^*}(-Z_g)) > P_g^{\pi^*}(-Z_g) - (P_g^\pi(-Z_g) - P_g^{\pi^*}(-Z_g)) = P_g^{\overline\pi^g}(-Z_g)$. Therefore $P_g^\pi(-Z_g) = P_g^{\underline\pi^g}(Z_g)$, and hence, by (13), $0.5 P_g^{\overline\pi^g}(-Z_g) + 0.5 P_g^{\underline\pi^g}(-Z_g) = P_g^{\pi^*}(-Z_g)$, which under mixture linearity implies $P_g^{0.5\overline\pi^g + 0.5\underline\pi^g}(-Z_g) = P_g^{\pi^*}(-Z_g)$. Parallel argument gives $0.5 P_f^{\overline\pi^f}(-Z_f) + 0.5 P_f^{\underline\pi^f}(-Z_f) = P_f^{0.5\overline\pi^f + 0.5\underline\pi^f}(-Z_f) = P_f^{\pi^*}(-Z_f)$.

Since there is a π-garbling from f to g, there must exist $\pi \in \Pi$ such that $P_f^\pi(-Z_g) \geq P_g^{\overline\pi^g}(-Z_g)$, and where $\pi' = \pi^* - (\pi - \pi^*) \in \Pi$, $P_f^{\pi'}(-Z_g) \leq P_g^{\underline\pi^g}(Z_g)$. As the π-garbling is centre-preserving, it must be that $0.5 P_f^\pi(-Z_g) + 0.5 P_f^{\pi'}(-Z_g) = P_f^{\pi^*}(-Z_g) = P_g^{\pi^*}(-Z_g) = 0.5 P_g^{\overline\pi^g}(-Z_g) + 0.5 P_g^{\underline\pi^g}(-Z_g)$. By mixture linearity we therefore have $P_g^{\hat\alpha\overline\pi^g + (1-\hat\alpha)\underline\pi^g}(-Z_g) \leq P_f^{\hat\alpha\pi + (1-\hat\alpha)\pi'}(-Z_g)$ if $\hat\alpha \geq 0.5$.

There are now two cases: first, if $P_g^{\hat\alpha\overline\pi^g + (1-\hat\alpha)\underline\pi_g}(-Z_g) = \theta$, mixture linearity and $P_g^{\hat\alpha\overline\pi^g + (1-\hat\alpha)\underline\pi^g}(-Z_g) \leq P_f^{\hat\alpha\pi + (1-\hat\alpha)\pi'}(-Z_g)$ imply that the capital holding required under book f must be at least as great as Z_g. Second, if $P_g^{\hat\alpha\overline\pi^g + (1-\hat\alpha)\underline\pi^g}(-Z_g) < \theta$, then for any $\epsilon > 0$ there must exist $\pi_1^g, \pi_2^g \in \Pi$ such that $P_g^{\hat\alpha\pi_1^g + (1-\hat\alpha)\pi_2^g}(-Z_g + \epsilon) > \theta$ – in which case parallel argument to that of the previous paragraph can show that there are $\pi_1^f, \pi_2^f \in \Pi$ such that $P_g^{\hat\alpha\pi_1^g + (1-\hat\alpha)\pi_2^g}(-Z_g + \epsilon) \leq P_f^{\hat\alpha\pi_1^f + (1-\hat\alpha)\pi_2^f}(-Z_g + \epsilon)$ if $\hat\alpha \geq 0.5$, and $\theta < P_g^{\hat\alpha\pi_1^g + (1-\hat\alpha)\pi_2^g}(-Z_g + \epsilon) \leq P_f^{\hat\alpha\pi_1^f + (1-\hat\alpha)\pi_2^f}(-Z_g + \epsilon)$ for all $\epsilon > 0$ similarly implies $Z_f \geq Z_g$. We therefore have $\hat\alpha \geq 0.5$ implies $Z_f \geq Z_g$.

To show the converse implication, we can reason in a parallel manner. As there is a π-garbling from f to g, it must be that where $\pi = \arg\max_{\pi \in \Pi} P_g^\pi(-Z_f)$ and $\pi' = \arg\min_{\pi \in \Pi} P_g^\pi(-Z_f)$, $P_g^\pi(-Z_f) \leq P_f^{\overline\pi_f}(Z_f)$, $P_g^{\pi'}(-Z_f) \geq P_f^{\underline\pi^f}(Z_f)$, and $0.5 P_g^\pi(-Z_f) + 0.5 P_g^{\pi'}(-Z_f) = P_g^{\pi^*}(-Z_f) = P_f^{\pi^*}(-Z_f) = 0.5 P_f^{\overline\pi^f}(-Z_f) + 0.5 P_f^{\underline\pi^f}(-Z_f)$. Thus we have $P_f^{\hat\alpha\overline\pi^f + (1-\hat\alpha)\underline\pi^f}(-Z_f) \leq P_g^{\hat\alpha\pi + (1-\hat\alpha)\pi'}(-Z_f)$ if $\hat\alpha \leq 0.5$. Reasoning as in the previous paragraph then shows that this implies $Z_g \geq Z_f$ in the two possible cases where $P_f^{\hat\alpha\overline\pi^f + (1-\hat\alpha)\underline\pi^f}(-Z_f) = \theta$ and where $P_f^{\hat\alpha\overline\pi^f + (1-\hat\alpha)\underline\pi^f}(-Z_f) < \theta$. Thus we have $\hat\alpha \leq 0.5$ implies $Z_f \leq Z_g$. $\square$

Proof of Proposition 4

Let $\overline\pi = \arg\max_{\pi \in \Pi} P_f^\pi(-Z_f)$ and $\underline\pi = \arg\min_{\pi \in \Pi} P_f^\pi(-Z_f)$ for book f and observe that $\overline\pi$ $(\underline\pi)$ must be the $\geq_U$-minimum (-maximum) of Π. By U-comonotonicity on $\{f, g\}$ it follows that $\overline\pi = \arg\max_{\pi \in \Pi} P_g^\pi(-Z_g)$ and $\underline\pi = \arg\min_{\pi \in \Pi} P_g^\pi(-Z_g)$ too. Following the proof of Proposition 2, it must

be that $P_g^{\hat{\alpha}\bar{\pi}+(1-\hat{\alpha})\underline{\pi}}(\{x : x \leq -Z_g\}) \geq \theta$. Using condition (3) of Proposition 3 we then have

$$P_g^{\hat{\alpha}\bar{\pi}+(1-\hat{\alpha})\underline{\pi}}(\{x : x \leq -Z_g\}) - P_g^{\hat{\alpha}'\bar{\pi}+(1-\hat{\alpha}')\underline{\pi}}(\{x : x \leq -Z_g\}) \leq$$
$$P_f^{\hat{\alpha}\bar{\pi}+(1-\hat{\alpha})\underline{\pi}}(\{x : x \leq -Z_f\}) - P_f^{\hat{\alpha}'\bar{\pi}+(1-\hat{\alpha}')\underline{\pi}}(\{x : x \leq -Z_f\}).$$

If $\hat{\alpha} = 0.5$, so that $\hat{\alpha}' > 0.5$, then

$$P_g^{\hat{\alpha}\bar{\pi}+(1-\hat{\alpha})\underline{\pi}}(\{x : x \leq -Z_g\}) = P_f^{\hat{\alpha}\bar{\pi}+(1-\hat{\alpha})\underline{\pi}}(\{x : x \leq -Z_f\}),$$

which implies that

$$P_f^{\hat{\alpha}'\bar{\pi}+(1-\hat{\alpha}')\underline{\pi}}(\{x : x \leq -Z_f\}) \geq P_g^{\hat{\alpha}'\bar{\pi}+(1-\hat{\alpha}')\underline{\pi}}(\{x : x \leq -Z_g\}) \geq \theta$$

and thus that $Z_f \geq Z_g$. Otherwise if $\hat{\alpha}' = 0.5$, so that $\hat{\alpha} < 0.5$, then

$$P_g^{\hat{\alpha}'\bar{\pi}+(1-\hat{\alpha}')\underline{\pi}}(\{x : x \leq -Z_g\}) = P_f^{\hat{\alpha}'\bar{\pi}+(1-\hat{\alpha}')\underline{\pi}}(\{x : x \leq -Z_f\}),$$

which implies that

$$P_g^{\hat{\alpha}\bar{\pi}+(1-\hat{\alpha})\underline{\pi}}(\{x : x \leq -Z_g\}) \geq P_f^{\hat{\alpha}\bar{\pi}+(1-\hat{\alpha})\underline{\pi}}(\{x : x \leq -Z_f\}) \geq \theta$$

and thus that $Z_g \geq Z_f$. $\square$