Centre for Climate Change Economics and Policy

The Munich Re Programme: *Evaluating the Economics of Climate Risks and Opportunities in the Insurance Sector*

Grantham Research Institute on Climate Change and the Environment

Ambiguity and insurance: robust capital requirements and premiums

Oliver Walker and Simon Dietz

November 2012

Centre for Climate Change Economics and Policy Working Paper No. 115

Munich Re Programme Technical Paper No. 16

Grantham Research Institute on Climate Change and the Environment

Working Paper No. 97

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 and has five inter-linked research programmes:

1. Developing climate science and economics
2. Climate change governance for a new global deal
3. Adaptation to climate change and human development
4. Governments, markets and climate change mitigation
5. The Munich Re Programme - Evaluating the economics of climate risks and opportunities in the insurance sector (funded by Munich Re)

More information about the Centre for Climate Change Economics and Policy can be found at: http://www.cccep.ac.uk.

The Munich Re Programme is evaluating the economics of climate risks and opportunities in the insurance sector. It is a comprehensive research programme that focuses on the assessment of the risks from climate change and on the appropriate responses, to inform decision-making in the private and public sectors. The programme is exploring, from a risk management perspective, the implications of climate change across the world, in terms of both physical impacts and regulatory responses. The programme draws on both science and economics, particularly in interpreting and applying climate and impact information in decision-making for both the short and long term. The programme is also identifying and developing approaches that enable the financial services industries to support effectively climate change adaptation and mitigation, through for example, providing catastrophe insurance against extreme weather events and innovative financial products for carbon markets. This programme is funded by Munich Re and benefits from research collaborations across the industry and public sectors.

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 in climate change and the environment. The Institute is funded by the Grantham Foundation for the Protection of the Environment and the Global Green Growth Institute, and has five research programmes:

1. Global response strategies
2. Green growth
3. Practical aspects of climate policy
4. Adaptation and development
5. Resource security

More information about the Grantham Research Institute on Climate Change and the Environment can be found at: http://www.lse.ac.uk/grantham.

Ambiguity and insurance: robust capital requirements and premiums

Oliver Walker[†] and Simon Dietz[*]

November 16, 2012

Abstract

Many insurance and reinsurance contracts are contingent on events such as hurricanes, terrorist attacks or political upheavals whose probabilities are not known with precision. There is a body of experimental evidence showing that higher premiums are charged for these "ambiguous" contracts, which may in turn inhibit (re)insurance transactions, but little research analysing explicitly how and why premiums are loaded in this way. In this paper we model the effect of ambiguity on the capital requirement of a (re)insurer whose objectives are profit maximisation and robustness. The latter objective means that it must hold enough capital to meet a survival constraint across a range of available estimates of the probability of ruin. We provide characterisations of when one book of insurance is more ambiguous than another and formally explore the circumstances in which a more ambiguous book requires at least as large a capital holding. This analysis allows us to derive several explicit formulae for the price of ambiguous insurance contracts, each of which identifies the extra ambiguity load.

JEL Classification Numbers: D81, G22.

Keywords: ambiguity, ambiguity aversion, ambiguity load, capital requirement, catastrophe risk, insolvency, insurance, more ambiguous, reinsurance, robustness, ruin, uncertainty, Solvency II.

[†]Department of Economics, University of Cambridge.
[*]Grantham Research Institute on Climate Change and the Environment, and Department of Geography and Environment, London School of Economics and Political Science.

We are very grateful to Pauline Barrieu and Trevor Maynard for comments on a previous draft. We would like to acknowledge the financial support of the Grantham Foundation for the Protection of the Environment, as well as the Centre for Climate Change Economics and Policy, which is funded by the UK's Economic and Social Research Council and by Munich Re. The usual disclaimer applies.

Email for correspondence: s.dietz@lse.ac.uk.
Telephone: +44 207 955 7589.
Fax: +44 207 107 5441.

1 Introduction

A variety of well-known principles exists for the technical pricing of insurance and reinsurance contracts when the probability of claims is known precisely. However, 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 "ambiguous". There may be several reasons why contracts are ambiguous, including a lack of historical, observational data, and the existence of competing theories, proferred by competing experts and formalised in competing forecasting models, of the casual processes governing events that determine their value. Ambiguity is thus 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 in the future.

There is by now a body of evidence to show that, faced with offering a contract under ambiguity, (re)insurers increase their premiums, limit coverage, or are unwilling to provide (re)insurance at all. In the academic literature, much of the evidence is survey-based: agents in the (re)insurance industry, including insurance actuaries and underwriters, and reinsurance underwriters, are asked to quote prices for hypothetical contracts in which the probabilities of loss are alternatively unknown or known (Cabantous, 2007; Hogarth and Kunreuther, 1989, 1992; Kunreuther et al., 1993, 1995; Kunreuther and Michel-Kerjan, 2009). Their responses reveal that prices for ambiguous contracts exceed prices for unambiguous contracts 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, there is also plenty of evidence that prices are increased or coverage limited in the presence of ambiguity. For example, it is common to find guidance to the effect that (re)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 directly explains 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 (re)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 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

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

II Directive (where it is called a Solvency Capital Requirement). Our twist is to require the capital held to be "robust", in the sense that the survival constraint is met no matter which of a set of estimates of the probability of ruin later turns out to be correct. Our focus is on so-called 'technical pricing', hence we abstract from competitiveness effects.

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 show conditions under which more ambiguous books entail higher capital holdings under our robust capital-setting rule. Since one book is more ambiguous than another only if any ambiguity-averse decision-maker would prefer it – and the (re)insurer would presumably prefer to hold as little capital as possible – we interpret our results as showing that our rule for robust capital-setting encodes a form of ambiguity aversion under the conditions we specify. We therefore demonstrate that our capital-setting rule is consistent with the behavioural evidence cited earlier.

A critical advantage of our capital-setting rule is that it allows us to derive pricing formulae for ambiguous contracts in a way that makes the effect of ambiguity on prices transparent. We therefore proceed, using a similar approach to Kreps (1990), to find pricing expressions under four different distributional assumptions about the (re)insurer's information. In each case, we identify an *ambiguity load* – distinct from the more familiar risk load – that is increasing in some measure of the ambiguity of the contract being priced. 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 be done using 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., 2010; Gollier, 2012). 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 assumes ambiguity-neutral insurers, because it is assumed that a process of Bayesian updating can collapse multiple probability measures (i.e. models) into a single posterior probability measure over claims/the returns on the book.

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 (re)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 underwriter's decision problem

We take the point of view of a (re)insurance underwriter who faces uncertainty over the performance of her book and wishes to maximise expected profits subject to a survival constraint. The classic treatment of this problem characterises the underwriter's uncertainty using a single probability measure over the space of events determining the book's return. Under this account the underwriter 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 underwriter's information does not take the form of a single probability measure over a space of relevant scenarios. In such cases, she 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 she faces. Such an underwriter is said to face *ambiguity*.

We model this kind of underwriter'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).

In the classic account of this problem, the underwriter is assumed to possess a single probability measure on $\mathcal{B}$, representing her information about payoff-relevant events. We, however, wish to allow for cases where the underwriter faces ambiguity and therefore endow her with a *set* of measures on $\mathcal{B}$, Π, encompassing all probability measures she believes might characterise her uncertainty correctly. We refer to Π as the set of *models*. Where the underwriter'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 $\mathcal{B}_\Pi$ representing the underwriter'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:

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

for any $E \in \mathcal{B}_\mathbb{R}$. In words, given the underwriter's beliefs about Π, $P_f(E)$ gives the probability the underwriter places on her 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. We write $P_f(\{y\})$ to indicate the probability under ν that f pays out precisely y.

Let us take as our starting point a familiar model of (re)insurance pricing based on maximising expected profit subject to a survival constraint (e.g. Kreps, 1990), in the tradition of Stone (1973).[2] In this model, there is a (re)insurer who, given any book $f \in \mathcal{F}$, sets her capital holding, Z_f, as follows:

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

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 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 (re)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 underwriter thus focuses on the single probability – as measured by P_f – of her book paying out less than her capital holding. In this sense she treats uncertainty as described within different models and uncertainty over which model is correct as equivalent.

We extend this framework to allow the underwriter to exhibit a further concern for *robustness*. Specifically, our paper concerns an underwriter who, given her uncertainty over the correct model governing her book's performance, wishes to limit the probability of ruin across *all models* in Π. Thus, where for any $\pi \in \Pi$ and $f \in \mathcal{F}$, the measure P_f^π on $\mathcal{B}_\mathbb{R}$ is defined as:

$$P_f^\pi(E) \quad = \quad \pi\left(f^{-1}(E)\right)$$

and we adopt the convention of using $P_f^\pi(y)$ for $P_f^\pi(\{x : x < y\})$, she sets Z_f according to:

$$Z_f \quad = \quad \min\{x : \max_{\pi \in \Pi} P_f^\pi(-x) \leq \theta\} \tag{2}$$

In contrast to (1), the capital-setting rule (2) requires the underwriter to limit the probability of overall ruin – as given by P_f^π – to no more than θ for all $\pi \in \Pi$. Z_f is thus determined by the "most adverse" model for the underwriter holding f in Π, and x is now the highest Value at Risk, with respect to θ, in Π. We shall only consider cases where (2) well defines Z_f. In Section 5 below we discuss the applicability of the rule (2) to insurance problems.

3 Ambiguity and the capital holding

Jewitt and Mukerji (2012) provide various choice-based accounts of what it is for one book 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 the 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. To characterise this we use

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

Kilbanoff et al.'s (2005, KMM) "smooth" representation of choice under ambiguity, under which preferences, given by the relation $\succeq$ over $\mathcal{F}$, are such that for any bounded $f, g \in \mathcal{F}$:

$$f \succeq g \iff \int_\Pi \phi\left(\int_S u(f(s))d\pi\right)d\nu \geq \int_\Pi \phi\left(\int_S u(g(s))d\pi\right)d\nu \tag{3}$$

where u and ϕ are continuous, strictly increasing functions, and ν and Π are the beliefs revealed by $\succeq$. Under this representation, the *ambiguity attitude* associated with $\succeq$ is given by the function ϕ: $\succeq$ is *ambiguity averse* if and only if the function ϕ is concave, ambiguity neutral iff ϕ is affine, and ambiguity seeking iff ϕ is convex. Where ϕ_O is Oskar's ambiguity attitude and ϕ_J is John's, Oskar is *more ambiguity averse* than John iff there exists a concave map ψ such that $\phi_O = \psi \circ \phi_J$.[3]

Note that (3) 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 KMM's representation, we would never be able to describe one unbounded book as being more or less ambiguous than another. And as we wish to do just this, we characterise "more ambiguous" relative to a narrower class of preferences than those consistent with KMM's representation. We thus define $\mathcal{P}_{\nu,\Pi}$ as the set of all preferences over $\mathcal{F}$ that are consistent with KMM's representation, that rank any $f, g \in \mathcal{F}$ according to (3) provided all the expectations in (3) are defined, and that share beliefs given by ν and Π. The ambiguity attitude of any $\succeq \in \mathcal{P}_{\nu,\Pi}$ is determined by the properties of the ϕ associated with $\succeq$ just as in KMM's representation. We say $\succeq \in \mathcal{P}_{\nu,\Pi}$ is *f-constrained* iff, given the functions u and ϕ associated with $\succeq$, $\int_\Pi \phi\left(\int_S u(f(s))d\pi\right)d\nu$ is defined.

The definition below provides the characterisation of "more ambiguous" that we use throughout what follows. We denote the symmetric component of $\succeq$ using $\sim$ as usual.

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

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

ii. For any f- and g-constrained $\succeq_A, \succeq_B \in \mathcal{P}_{\nu,\Pi}$ 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 "$\mathcal{P}_{\nu,\Pi}$-" and simply say f is "more ambiguous" than g. Note condition (i) above and (3) imply that if f is more ambiguous than g, then $P_f = P_g$. 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. Under the assumption

[3] The characterisations of ambiguity aversion, ambiguity neutrality, and "more ambiguity averse" are all founded on more primitive definitions, which are equivalent to the properties of ϕ under KMM's representation. See Jewitt and Mukerji (2012) for details.

that supp(ν) = Π, 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 $\pi' \in \Pi$:

$$P_g^{\pi'}(E) \quad = \quad \int_\Pi P_f^{\pi}(E) dK_{\pi'} \tag{4}$$

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.

Jewitt and Mukerji describe a Markov kernel K from $(\Pi, \mathcal{B}_\Pi)$ to itself as *ν-preserving* if and only if, for all $E \in \mathcal{B}_\Pi$,

$$\mu(E) \quad = \quad \int_\Pi K_\pi(E) d\nu.$$

Whenever there is a ν-preserving π-garbling from f into g, $P_f(E) = P_g(E)$ for all $E \in \mathcal{B}_\mathbb{R}$. Thus, where this is so, the underwriter believes the likelihood of f and g paying out any amount is the same, but she expects her judgement of this likelihood to change more for f than for g upon learning the true model in Π. Intuitively, a ν-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 ν-preserving π-garbling preserves a book's *expected payoff-distribution*, but makes this 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 ν-preserving π-garbling from f to g then f is more ambiguous than g.*

It thus follows that if there is a ν-preserving π-garbling from f into g, then any ambiguity-averse agent whose preferences belong to $\mathcal{P}_{\nu,\Pi}$ would prefer f to g. Our treatment of the (re)insurer's problem does not specify her preferences – we merely require her to set her capital requirement according to the rule (2) – so we will not be able to gauge her her ambiguity attitude explicitly in the sense of KMM's representation. Nonetheless, if we suppose that the (re)insurer's shareholders are able to diversify whatever ambiguity they face, then since a higher capital holding implies a greater cost to the insurer, it seems natural to suppose that if $P_f = P_g$ but $Z_f > Z_g$, the underwriter would prefer to hold g than f. Thus, if we show that f's being more ambiguous than g implies that f incurs a higher capital requirement, we could conclude that the decision-making rule (2) is ambiguity averse. This is in contrast to decision rule (1), which sets $Z_f = Z_g$ whenever f is more ambiguous than g and can thus be regarded as ambiguity neutral.

Our first result shows precisely this, given the sufficient condition from Proposition 1 for f to be more ambiguous than g.

Proposition 2 *Suppose Z_f and Z_g are well defined by (2). Then if there is a ν-preserving π-garbling from f to g, $Z_f \geq Z_g$.*

Proof: Suppose there is a π-garbling from f to g and let $\pi_g^* = \arg\max_{\pi \in \Pi} P_g^{\pi}(-Z_g)$. By (2) it must be that $P_g^{\pi_g^*}(-Z_g) \leq \theta$, and if $P_g^{\pi_g^*}(-Z_g) < \theta$ then $P_g^{\pi_g^*}(-Z_g) + P_g^{\pi_g^*}(\{-Z_g\}) \geq \theta$. Thus, $P_g^{\pi_g^*}(\{x : x \leq -Z_g\}) \geq \theta$. By (4), $\max_{\pi \in \Pi} P_f^{\pi}(\{x : x \leq -Z_g\}) \geq P_g^{\pi}(\{x : x \leq -Z_g\})$ for all $\pi \in \Pi$ and hence $\max_{\pi \in \Pi} P_f^{\pi}(\{x : x \leq -Z_g\}) \geq P_g^{\pi_g^*}(\{x : x \leq -Z_g\}) \geq \theta$. This implies $Z_f \geq Z_g$. $\square$

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. U-comonotonicity is likely to hold in many applications – indeed some of the special cases we examine in Section 4 below require it – so it is of interest to show that (2) is an ambiguity-averse rule for all books under the condition. This is the content of our next result.

To begin with we define U-comonotonicity.[4]

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

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

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 $\pi \geq_U \pi'$ iff π places a lower probability on the extreme weather event than π' does.

To state Jewitt and Mukerji's characterisation of "more ambiguous" under U-comonotonicity, we need some additional notation. Let X_L be the collection of all lower intervals in $\mathbb{R}$ and – assuming Π is U-comonotone for some set of books – use Π_L for the

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

collection of all $\geq_U$-lower intervals in Π. That is:

$$\begin{aligned} X_L &= \{\{x : x \leq x'\} : x' \in \mathbb{R}\} \cup \{\{x : x < x'\} : x' \in \mathbb{R}\} \\ \Pi_L &= \{\{\pi : \pi \leq_U \pi'\} : \pi' \in \Pi\} \cup \{\{\pi : \pi <_U \pi'\} : \pi' \in \Pi\} \end{aligned}$$

Now define $P_{f,\nu}$ as the probability measure on $\mathcal{B}_{\mathbb{R}} \times \mathcal{B}_{\Pi}$ that satisfies:

$$P_{f,\nu}(E \times F) = \int_F P^f_\pi(E) d\nu$$

for any $E \in \mathcal{B}_{\mathbb{R}}$ and $F \in \mathcal{B}_{\Pi}$. That is, given the underwriter's beliefs, $P_{f,\nu}(E \times F)$ is the probability that the true measure lies in F and f pays out in E.

Jewitt and Mukerji's characterisation[5] ambiguity aversion under U-comonotonicity is as follows.

Proposition 3 *[Jewitt-Mukerji 2] Suppose* Π *is U-comonotone on* $\{f, g\}$*. Then the following two statements are equivalent:*

1. *f is more ambiguous than g.*
2. *For all $E \in \mathcal{B}_{\mathbb{R}}$, $P_f(E) = P_g(E)$ and for any $E \times F \in X_L \times \Pi_L$, $P_{f,\nu}(E \times F) \geq P_{g,\nu}(E \times F)$.*

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 and thus the probability of any adverse payoff (any E in X_L) conditional on an adverse model (any F in Π_L) is higher under f than under g.

The result allows us to show that the capital-holding rule (2) is ambiguity averse for any pair of acts, f and g provided Π is U-comonotone for $\{f, g\}$.

Proposition 4 *If* Π *is U-comonotone on* $\{f, g\}$ *and* f *is more ambiguous than* g *then* $Z_f \geq Z_g$*.*

Proof: Let $\pi^* = \arg\max_{\pi \in \Pi} P^\pi_f(\{x : x \leq -Z_f\})$ and observe that π^* must be the $\geq_U$-minimum of Π. Thus by U-comonotonicity it follows that $\pi^* = \arg\max_{\pi \in \Pi} P^\pi_g(\{x : x \leq -Z_g\})$. As in the proof of Proposition 2 it must be that $P^{\pi^*}_g(\{x : x \leq -Z_g\}) \geq \theta$. Since $\{x : x < -Z_g\} \times \{\pi^*\} \in X_L \times \Pi_L$, Proposition 3 implies $P^{\pi^*}_f(\{x : x \leq -Z_g\}) = P_{f,\mu}(\{x : x \leq -Z_g\} \times \{\pi^*\}) \geq P_{g,\mu}(\{x : x \leq -Z_g\} \times \{\pi^*\}) = P^{\pi^*}_g(\{x : x \leq -Z_g\}) \geq \theta$, and hence that $Z_f \geq Z_g$. $\square$

[5]The result reported here is slightly different to that in Jewitt and Mukerji, who define U-comonotonicity in terms of 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.

4 Contract Pricing under Ambiguity

We now examine the impact of ambiguity on a more practical level, deriving explicit expressions for the price of an individual contract given competitive behaviour and a capital holding set according to (2). In order to obtain expressions that clearly show how introducing ambiguity leads to a departure from the benchmark pricing formula in the absence of ambiguity, we need to make assumptions about the distribution of payoffs under each model and the distribution of model parameters under the measure ν. We examine four cases, each of which nests the benchmark analysis, chosen on the basis of their analytical tractability and applicability to real insurance problems. We only consider contract pricing, thus we ignore deductibles, co-insurance and other design options that an insurer might use.

4.1 Benchmark: no ambiguity

As a benchmark for what follows and in order to introduce terminology, we begin by reviewing the case examined by Kreps (1990) where the underwriter's information is unambiguous and payoffs are normally distributed. 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 $\{S, \mathcal{B}\}$ is normal for all $\pi \in \Pi$ with 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's worth emphasising that in this framework we assume that the underwriter sets her capital holding according to rule (2) – that is, she does exhibit a concern for robustness as outlined in the Introduction – and we allow Π to be non-singleton – implying that, across the class of all books, the underwriter may face some ambiguity. However, because we restrict our focus to $\mathcal{F}_0$, a class of unambiguous books, the underwriter faces no ambiguity and therefore her capital holding rule is equivalent to that in (1). We adopt this approach in order to make clearer the generalisations to richer sets of books in subsequent sections.

Where Φ is the cdf of a standard normal and $-z = \Phi^{-1}(\theta)$, the underwriter's capital holding for $f \in \mathcal{F}_0$ is determined by:

$$Z_f \;=\; z\sigma_f - \mu_f \tag{5}$$

We consider an insurer endowed with book f who agrees to an additional contract c, which is itself a book in $\mathcal{F}_0$, thereby ending up with book $f' = f + c$. As a result of signing c, she needs to increase her 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:

$$\mu_c \;=\; y\left(Z_{f'} - Z_f\right) \tag{6}$$

Given (5) 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:

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

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

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

Under assumptions described by Kreps, $\mathcal{R}_{c,f}$ may be approximated by $(yz/(1+y))\,(\rho_{c,f} + \sigma_c/2\sigma_f)$.

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 of c (including administrative costs) to the insurer, L_c, and we say an insurer is *competitive*[6] whenever it always sets P_c such that μ_c satisfies (6). We can now state Kreps's pricing result, the proof of which is immediate from the analysis above.

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

$$P_c \quad = \quad L_c + \mathcal{R}_{c,f}\sigma_c \tag{7}$$

Kreps calls the expression $\mathcal{R}_{c,f}\sigma_c$ the *risk load* for contract c. Note that it arises solely as a consequence of the underwriter'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 its actuarially fair cost, 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 relative riskiness of the contract compared to the pre-existing book (σ_c/σ_f), and 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; standard deviation known

4.2.1 Mean uniformly distributed

Our first generalisation of Kreps's framework considers 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 $\pi \in \Pi$, $f(s)$ is normal under P_f^π on $\{S, \mathcal{B}\}$ with 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

[6]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.

[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.

(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 $f, f' \in \mathcal{F}_1$:

$$\mu^{\pi}_{f'} = a_{f'} + (\mu^{\pi}_f - a_f)\frac{(b_{f'} - a_{f'})}{(b_f - a_f)} \tag{8}$$

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 underwriter has a collection of forecasts at her 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, she is sure that the variance is as reported – σ^2_f for book f – and thinks the true mean must lie somewhere between these optimistic and pessimistic bounds.

She constructs Π and ν using three assumptions.

I The members of Π are ordered according to their pessimism so that (8) is satisfied and for any $c \in [a_f, b_f]$, $\mu^{\pi}_f = c$ for one $\pi \in \Pi$;

II The distribution of each book f's payoffs are approximated as normal with mean μ^{π}_f and variance σ^2_f under each measure P^{π}_f corresponding to a $\pi \in \Pi$; and

III ν is set such that condition (1.ii), imposing a uniform distribution on μ^{π}_f, holds.

I may be justified in case the underwriter finds it reasonable while II is a standard procedure in empirical modelling. Assumption III is reasonable provided she has no evidence to suggest any value of μ^{π}_f in $[a_f, b_f]$ is more plausible than any other, in which case the uniformity of μ^{π}_f follows from the principle of insufficient reason.

Under these assumptions, any book she considers belongs to $\mathcal{F}_1$ given Π and ν.

Given the decision rule (2), for any $f \in \mathcal{F}_1$ we have:

$$Z_f \quad = \quad z\sigma_f + 3\text{Var}[\mu^{\pi}_f] - \mu_f \tag{9}$$

where $\text{Var}[\mu^{\pi}_f]$ is the variance of the random variable μ^{π}_f (equal to $(1/12)(b_f - a_f)^2$ 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 (9) and (6) as above, we obtain:

$$\begin{aligned} \mu_c &= \mathcal{R}_{c,f}\sigma_c + \frac{3y}{1+y}\left(\text{Var}[\mu^\pi_{f'}] - \text{Var}[\mu^\pi_f]\right) \\ &= \mathcal{R}_{c,f}\sigma_c + \frac{3y}{1+y}\left(\text{Var}[\mu^\pi_c] + 2\text{Cov}[\mu^\pi_c, \mu^\pi_f]\right) \end{aligned}$$

Using the fact that, given (8), $\text{Cov}[\mu^\pi_c, \mu^\pi_f] = \text{sd}[\mu^\pi_f]\text{sd}[\mu^\pi_c]$ and defining

$$\mathcal{A}_{c,f,1} := \left(\frac{3y}{1+y}\right)\left(1 + \frac{2\text{sd}[\mu^\pi_f]}{\text{sd}[\mu^\pi_c]}\right)$$

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:*

$$P_c = L_c + \mathcal{R}_{c,f}\sigma_c + \mathcal{A}_{c,f,1}\text{Var}[\mu^\pi_c]$$

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 $\text{Var}[\mu^\pi_c] = 0$ and so P_c is set according to (7). However, if c is in $\mathcal{F}_1 \setminus \mathcal{F}_0$ – that is to say c is ambiguous – then P_c also incorporates an *ambiguity load* equal to $\mathcal{A}_{c,f,1}\text{Var}[\mu^\pi_c]$. The ambiguity load arises because c is ambiguous and the underwriter's decision rule encodes a concern for robustness. It is increasing in $\text{Var}[\mu^\pi_c]$, which may be thought of as an approximate measure of how ambiguous c is[8], implying that insurers facing this structure of uncertainty charge higher premia for more ambiguous contracts, and it is increasing in the ambiguity of the pre-existing book (measured by $\text{sd}[\mu^\pi_f]$).

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 $\pi \in \Pi$, $f(s)$ is normal under P^π_f on $\{S, \mathcal{B}\}$ with mean μ^π_f and variance σ^2_f; (2.ii) for all $f \in \mathcal{F}_2$, μ^π_f has a symmetric triangular distribution on $[a_f, b_f]$ given ν on $\{\Pi, \mathcal{B}_\Pi\}$; (2.iii) $\mathcal{F}_2$ is closed under addition; 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 (8) 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 underwriter from Example 1 thinks that, for any book f, values of μ^π_f 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^π_f is triangularly

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

distributed[9] with minimum a_f, maximum b_f, and mode $(a_f+b_f)/2$, she might proceed using assumptions I and II as above, but changing III to make sure that ν is such that (2.ii) is satisfied. Given Π and ν thus constructed, every book she considers will belong to $\mathcal{F}_2$.

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

$$Z_f \quad = \quad z\sigma_f + \sqrt{6}\text{sd}[\mu_f^\pi] - \mu_f$$

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

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

Where we use $\mathcal{A}_{c,f,2}$ to denote $(\sqrt{6}y/(1+y))(2\text{sd}[\mu_f^\pi] + \text{sd}[\mu_c^\pi])/(\text{sd}[\mu_f^\pi] + \text{sd}[\mu_{f'}^\pi])$, which takes the approximate value of $(\sqrt{6}y/(1+y))(1+\text{sd}[\mu_c^\pi]/2\text{sd}[\mu_f^\pi])$ under Kreps's assumptions, this gives us:

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

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:*

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

Once again, the result generalises Proposition 5 by incorporating an ambiguity load that is zero for $c \in \mathcal{F}_0$ and increasing in a measure of the ambiguity of c, $\text{sd}[\mu_c^\pi]$. By contrast to Proposition 6, however, Proposition 7 implies that P_c is decreasing in the ambiguity of f (as measured by $\text{sd}[\mu_f^\pi]$), which means that insurers with ambiguous books should be more able to accommodate ambiguous contracts.

4.3 Mean known; standard deviation 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 $\pi \in \Pi$, $f(s)$ is normal under P_f^π on $\{S, \mathcal{B}\}$ with mean μ_f and variance $\left(\sigma_f^\pi\right)^2$; (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:

$$\sigma_{f'}^\pi \quad = \quad a_{f'} + (\sigma_f^\pi - a_f)\frac{(b_{f'} - a_{f'})}{(a_f - b_f)} \qquad (10)$$

[9]We choose this distribution as (2) does not well-define the capital holding unless μ_f^π has a bounded support.

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

We imagine this case applying to an underwriter 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.[10]

Working as before, we have:

$$Z_f \quad = \quad z(\mathbb{E}[\sigma_f^\pi] + 3\text{Var}[\sigma_f^\pi]) - \mu_f \tag{11}$$

for any $f \in \mathcal{F}_3$. And thus, where $f, c \in \mathcal{F}_3$ and $f' = f + c$, (6) implies:

$$\mu_c \quad = \quad \frac{yz}{1+y}\left((\mathbb{E}[\sigma_{f'}^\pi] - \mathbb{E}[\sigma_f^\pi]) + 3(\text{Var}[\sigma_{f'}^\pi] - \text{Var}[\sigma_f^\pi])\right)$$

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 obtain

$$\begin{aligned}
\mu_c \quad &= \quad \mathbb{E}[\mathcal{R}_{c,f}\sigma_c^\pi] + \frac{3yz}{1+y}(\text{Var}[\sigma_{f'}^\pi] - \text{Var}[\sigma_f^\pi]) \\
&= \quad \mathbb{E}[\sigma_c^\pi]\mathbb{E}[\mathcal{R}_{c,f}] + \text{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}] + \frac{3yz}{1+y}(\text{Var}[\sigma_{f'}^\pi] - \text{Var}[\sigma_f^\pi])
\end{aligned}$$

Now note that $\text{Var}[\sigma_{f'}^\pi] - \text{Var}[\sigma_f^\pi] = \text{Var}[\sigma_{f'}^\pi - \sigma_f^\pi] - 2\text{Var}[\sigma_f^\pi] + 2\text{Cov}[\sigma_{f'}^\pi, \sigma_f^\pi]$ and, given (10), one may verify that $\text{Var}[\sigma_{f'}^\pi - \sigma_f^\pi] = \text{Var}[\sigma_c^\pi]$ and $\text{Cov}[\sigma_{f'}^\pi, \sigma_f^\pi] = \text{Var}[\sigma_f^\pi] + \text{sd}[\sigma_f^\pi]\text{sd}[\sigma_c^\pi]$. Thus, defining $\mathcal{A}_{c,f,3} = \frac{3yz}{1+y}\left(1 + \frac{\text{sd}[\sigma_f^\pi]}{\text{sd}[\sigma_c^\pi]}\right)$, we have:

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

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:*

$$P_c \quad = \quad L_c + \mathbb{E}[\sigma_c^\pi]\mathbb{E}[\mathcal{R}_{c,f}] + \text{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}] + \mathcal{A}_{c,f,3}\text{Var}[\sigma_c^\pi]$$

Once again, whenever $c \in \mathcal{F}_0$, the pricing formula above reduces to (7). 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, $\mathcal{A}_{c,f,3}\text{Var}[\sigma_c^2]$, that is increasing in the ambiguity of c. Like the corresponding term in Proposition 6 but unlike its counterpart in Proposition 7, it is also increasing in the ambiguity of the pre-existing book f. 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. Note that the presence of this extra term implies that, in contrast to the other cases examined so far, the ambiguity load for $f, c \in \mathcal{F}_3$ could be negative.

[10]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 underwriter could equally invoke the principle to impose the uniformity of $\left(\sigma_f^\pi\right)^2$, in which case the collection of books she considers could not satisfy (3.ii).

4.4 Mean and Standard Deviation Uncertain

As a final exercise, we consider an informational structure that nests two of the cases described above: where both the mean and the standard deviation 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 $\pi \in \Pi$, $f(s)$ is normal under P_f^π on $\{S, \mathcal{B}\}$ with 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 (8) and (10) (the latter with obvious relabelling). Apart from cases where $\mathcal{F}_4 = \mathcal{F}_0$, Π is not U-comonotone for $\mathcal{F}_4$.

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

$$Z_f \quad = \quad z\left(\mathbb{E}[\sigma_f^\pi] + 3\mathrm{Var}[\sigma_f^\pi]\right) + 3\mathrm{Var}[\mu_f^\pi] - \mu_f$$

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

$$\mu_c \quad = \quad \frac{yz}{1+y}\left((\mathbb{E}[\sigma_{f'}^\pi] - \mathbb{E}[\sigma_f^\pi]) + 3(\mathrm{Var}[\sigma_{f'}^\pi] - \mathrm{Var}[\sigma_f^\pi])\right) + \frac{3y}{1+y}\left(\mathrm{Var}[\mu_{f'}^\pi] - \mathrm{Var}[\mu_f^\pi]\right)$$

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:*

$$P_c \quad = \quad L_c + \mathbb{E}[\sigma_c^\pi]\mathbb{E}[\mathcal{R}_{c,f}] + \mathrm{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}] + \mathcal{A}_{c,f,1}\mathrm{Var}[\mu_c^\pi] + \mathcal{A}_{c,f,3}\mathrm{Var}[\sigma_c^\pi]$$

Thus, where books and contracts belong to $\mathcal{F}_4$, the ambiguity load for any contract is simply the sum of a component ($\mathcal{A}_{c,f,1}\mathrm{Var}[\mu_c^\pi]$) arising due to ambiguity in the contract's mean and a component ($\mathrm{Cov}[\sigma_c^\pi, \mathcal{R}_{c,f}] + \mathcal{A}_{c,f,3}\mathrm{Var}[\sigma_c^\pi]$) reflecting ambiguity in its standard deviation. This additive structure 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 insurance under a robust capital-setting rule. We show, at a general level, that under our capital-setting rule, increasing ambiguity leads to higher capital-holdings and thus to higher costs. 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.

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[11]. Furthermore, in some regions the regulatory framework governing insurers' capital holding approximates our rule to some degree. For example, the EU's Solvency II Directive mandates (via the Solvency Capital Requirement) a capital holding to limit the probability of ruin (θ in (2)) at 0.005 over a one-year horizon and explicitly demands that insurer valuation models fulfil robustness conditions (Lloyd's, 2010). In these cases, (2) has obvious descriptive applicability.

Considering the rule from a normative perspective, we note that the concern for robustness encoded in (2) takes an analogous (if not identical) form to the decision rule in Hansen and Sargent (2008). The latter has been shown to be equivalent to Gilboa and Schmeidler's (1989) axiomatically founded "maxmin" expected utility model in (Hansen and Sargent, 2001). However, these decision rules 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 their ambiguity (as measured by the variance of their uncertain distributional parameters). 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 (re)insurer to "hedge" against the ambiguity in her 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.

Another interesting feature of our contract-pricing results is that they show no clear relation between the ambiguity of the pre-existing book and the price of a contract: for example, Proposition 6, where the mean payoff is uniform, makes P_c increasing in the ambiguity of the insurer's book, while Proposition 7, where the mean payoff is triangular, states the opposite. Which, if either, of these claims is correct tells us whether, in an optimally structured insurance industry, some firms should "specialise" in ambiguous insurance contracts while others avoid them or all firms should bear the same degree of ambiguity. It is thus of interest that our capital-holding rule does not, in itself, give a definitive answer to this question.

[11] 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 at odds with 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 (re)insurer's ruin, as in our framework. Indeed, the risk of ruin is known to be an important consideration more generally when (re)insurers set capital holdings and price contracts, especially for catastrophe risks (e.g. Kunreuther and Michel-Kerjan, 2009).

References

Alary, D., C. Gollier, and N. Treich (2010): "The effect of ambiguity aversion on insurance demand," Toulouse School of Economics.

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.

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," *The Quarterly Journal of Economics*, 75, 643–669.

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.

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

Gollier, C. (2012): "Optimal insurance design of ambiguous risks," Institut d'Économie Industrielle (IDEI), Toulouse, Working Paper Series No. 718.

Hansen, L. and T. Sargent (2001): "Robust control and model uncertainty," *The American Economic Review*, 91, 60–66.

——— (2008): *Robustness*, Princeton Univ Pr.

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," *The 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.

Lloyd's (2010): "Solvency II: Detailed guidance notes for dry run process," .

Maccheroni, F., M. Marinacci, and D. Ruffino (2010): "Alpha as ambiguity: robust mean-variance portfolio analysis," *IGIER Working Paper 373, Bocconi University.*

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.