DISAGREEMENT AVERSION

Antoine Bommier*, Adrien Fabre*, Arnaud Goussebaïle* and Daniel Heyen†‡

March 18, 2022

Link to latest version

Abstract

Experts often disagree. A decision-maker may be averse to such expert disagreement. Existing models of aversion to expert disagreement rest on ambiguity-averse preferences adopting a unanimity principle: If all experts consider one choice better than another, so should the decision-maker. Such unanimity among experts, however, can be spurious, masking substantial disagreement on the underlying reasons. We introduce a novel notion of disagreement aversion to distinguish spurious from genuine unanimity and develop a model that can capture disagreement aversion in our sense. The central element of our model is the cautious aggregation of experts' beliefs.

JEL: D81; D83; D71.

Keywords: Disagreement Aversion; Ambiguity Aversion; Belief Aggregation; Decision under Uncertainty; Precautionary Principle.

*ETH Zürich. Bommier: abommier@ethz.ch (corresponding author).

†University of Kaiserslautern and ETH Zürich.

‡We are grateful to Loïc Berger, Adam Dominiak, Hans Gersbach, Christoph Rheinberger, Lorenzo Stanca, Jean-Marc Tallon, Massimo Tavoni and Peter Wakker for their helpful and generous comments. We also thank the seminar participants at ETH Zürich, Global Priorities Institute at Oxford, HEC Lausanne, KU Leuven, Paris School of Economics and Toulouse School of Economics. We are grateful to Rémi Peyre for sharing mathematical insights. We thank Carmel Connell for the proofreading.

Experts disagree. Such expert disagreement complicates decisions on topics as diverse as climate change, macroeconomics, and the Covid-19 pandemic. The different, potentially conflicting scientific opinions force a decision-maker to act without precise knowledge of the likelihoods of the possible outcomes. In such a situation, the decision-maker may conclude that there is scientific uncertainty and that most likely none of the experts is completely right. Accordingly, she may be averse to expert disagreement and prefer "safer" actions on which experts differ less. This paper introduces a novel notion of aversion to expert disagreement and illustrates its implications.

To fix ideas, consider the following simple example as summarized in Table 1. The CEO of a firm needs to decide whether to implement a project or not. The project's net return depends on contingencies over which the CEO has no control: The net return is $+10$ in the "good" state, 0 in the "neutral" state, and -10 in the "bad" state (we assume that the CEO is risk-neutral, or that the numbers in Table 1 are utils). The CEO consults two experts for their views about the likelihood of each state. Whereas the first expert is convinced that the neutral state will occur for sure, the second expert assigns a null probability to this event and considers that the good and bad states are equally likely.

Table 1 – Leading example of a CEO's decision whether to implement a project.

	Payoffs (in utils)		Beliefs (probabilities)	
	No project	Project	Expert 1	Expert 2
State 1 ("bad")	0	-10	0	0.5
State 2 ("neutral")	0	0	1	0
State 3 ("good")	0	10	0	0.5

Say both experts share the CEO's objective to maximize expected return (or utility). Then, both experts agree that implementing the project is as good as not implementing it, since both options provide the same expected utility. However, the experts provide very different reasons for this conclusion. A cautious CEO may consider the possibility that the expected return of the project is worse than what any expert foresees: Maybe expert 1 is right that implementing the project cannot yield positive profits, and expert 2 is right that losses are a likely outcome.

Intuitively, the CEO may infer from the underlying expert disagreement (despite agreement on the expected utility) that there is no robust understanding of what the project's outcomes could be and consequently prefer the no-project option about which experts (fully) agree.

There are several papers in the decision theory and applied literature on how to make decisions when facing imprecise or conflicting information (Baillon, Cabantous and Wakker 2012; Berger et al. 2021; Crès, Gilboa and Vieille 2011; Gajdos and Vergnaud 2013; Gajdos et al. 2008; Hill 2012; Mongin 1995, 1998; Nascimento 2012; Stanca 2021). All these papers keep as a fundamental assumption a unanimity principle[1]: If according to all experts implementing the project is just as good as doing nothing, the decision-maker should also be indifferent between the two options. That experts hold diverging beliefs has no impact on the decision-maker's willingness to proceed with the project. A few papers have criticized the unanimity principle for preference aggregation under uncertainty, with Gilboa, Samet and Schmeidler (2004) and Mongin (2016) as prominent examples.[2] Both argue against the unanimity principle (or Pareto condition) when people hold different beliefs and tastes. Gilboa, Samet and Schmeidler (2004) give the example of a duel: Both gentlemen agree to fight because each one wants to win (conflicting tastes) and each one thinks he will win (contradictory beliefs). Despite the unanimity among duelists, the authors argue that the duel should be prevented because "Society should not necessarily endorse a unanimous choice when it is based on contradictory beliefs", as these can spuriously reconcile conflicting tastes. Mongin (2016) argues further that the unanimity of individuals' opinions does not have to be reflected at the social level whenever the unanimity was obtained despite disagreement on what individuals use to form their opinions. Divergence in beliefs could then be a sufficient reason to depart from the unanimity principle even if people share the same objective. In our introductory example, there is no conflict on tastes, as all experts agree that the objective is to maximize expected net return. There is spurious unanimity in the sense that experts disagree on the probability of relevant states, but

[1]Various terminology is used to refer to this unanimity principle. For example, Gajdos et al. (2008) call it "dominance", while it is named "unanimity" in Crès, Gilboa and Vieille (2011) and Stanca (2021) refers to it as the "Pareto condition".

[2]Other examples include Machina and Siniscalchi (2014) and Skiadas (2013).

happen to agree on the preference ordering of the project and no-project options. In line with Mongin (2016), we argue that the CEO does not have to adopt the experts' unanimous conclusion. In other words, we highlight that spurious unanimity can occur even when tastes are commonly shared, and apply Mongin's general criticism of the unanimity principle to such situations.

This paper makes three contributions. The first is to formalize a novel notion of disagreement aversion that has traction even in the presence of spurious unanimity resulting exclusively from diverging beliefs, as in our introductory example. Intuitively, disagreement aversion is a tendency to prefer choices on which experts have reached a consensus. However, what is meant by "consensus" merits further consideration. One possible interpretation of consensus on a choice option is that all experts agree on the resulting expected utility. However, as explained above, such *utility-consensus* may be "spurious" and obtained despite fundamental heterogeneity in expert beliefs. For this reason, we introduce a stronger notion of consensus, *distribution-consensus*, which requires that all experts not only agree on the expected utility but even hold a consensus on the distribution of outcomes. The different notions of aversion to experts disagreement defined in the literature, such as "ambiguity aversion" (Ghirardato and Marinacci 2002) or "imprecision aversion" (Gajdos et al. 2008), all rest on the unanimity principle and are thus intimately tied to the concept of utility-consensus. The existing notions of disagreement aversion, therefore, do not have bite if experts have diverging beliefs but happen to agree on the expected utility of a choice option. Our novel notion of disagreement aversion, in contrast, captures a decision-maker's aversion to a lack of distribution-consensus. Going back to the introductory example, both the "no-project" and the "project" options are utility-consensual; therefore, under any model that fulfills the unanimity principle, the decision-maker is necessarily indifferent between both options. In contrast, only the "no-project" option is distribution-consensual. Accordingly, as experts have no distribution-consensus on the project option, a decision-maker averse to disagreement (in our sense) prefers to abstain from the project.

Our second contribution is to introduce a model that allows for disagreement aversion. Like most previous contributions, we borrow from the ambiguity aversion literature. However, since disagreement aversion requires abandoning the unanimity principle, we cannot rely on the most common ambiguity models. Indeed, those "monotone models" assume a property of monotonicity (in terms of states) which in our setting translates to the unanimity principle (in terms of experts). We instead build upon the dual model of ambiguity aversion by Bommier (2017) and introduce a "distribution-averaging model". In such a model, for each choice option separately, the decision-maker aggregates the distribution functions of outcomes provided by the experts and uses that aggregated distribution function to evaluate the choice option. Importantly, the aggregation is specific to the choice option, leaving the possibility of a cautious aggregation that gives greater weight to more pessimistic views. Formally, this aggregated distribution function is obtained by aggregating pointwise the decumulative distribution functions provided by all experts into a single decumulative function.[3] Intuitive characteristics of the aggregator (or "averager") such as concavity ensure that the decision-rule is disagreement-averse according to our novel notion. We contrast our distribution-averaging model with the approach based on monotone ambiguity models. Both approaches rely on some form of cautious aggregation where the degree of cautiousness is what drives aversion to expert disagreement. The fundamental difference is that in our model the aggregation occurs at the level of experts' beliefs, while in monotone ambiguity models the aggregation occurs at the level of expected utility values, yielding what we call "expected-utility-averaging models". Although structurally different, these approaches can be compared in terms of ambiguity aversion. In particular, once utility is normalized to take values in $[0,1]$, like probabilities, we show that for a given strictly concave averager and given risk preferences, our disagreement-averse specification exhibits greater ambiguity aversion than the corresponding expected-utility-averaging model. In that sense, our paper brings the notion of aversion to expert disagreement one step further.

The third contribution is to highlight the implications of our novel notion of disagreement aversion in concrete applications. We show that greater disagreement aversion always results in more cautious choices, i.e. choices that reduce the dis-

[3] If F denotes the cumulative distribution function of a continuous random variable, $1-F$ is the decumulative distribution function. We work with decumulative (rather than cumulative) distribution functions to obtain the result that disagreement aversion is implied by a concave (rather than convex) aggregator.

persion of ex-post utility, such as higher climate mitigation or higher precautionary savings. We compare our findings to those obtained with expected-utility-averaging models. When experts' beliefs can be unambiguously ranked from the most optimistic to the least optimistic (in terms of first-order stochastic dominance), then more ambiguity aversion also leads to more cautious choices in expected-utility-averaging models. The result, however, does not extend to the case where experts' beliefs cannot be ranked in terms of first-order stochastic dominance. We give two examples – on climate mitigation and precautionary savings, respectively – where an increase in ambiguity aversion in expected-utility-averaging models leads to *less cautious* choices, namely lower mitigation and lower savings. The effects are thus in direct opposition to what we obtain with distribution-averaging models for which an increase in ambiguity (or disagreement) aversion leads to more cautious choices.

Our paper is related to several strands of the literature. There is a vast corpus of works in mathematics and management on *opinion pooling*, in particular linear opinion pools in which the decision-maker uses a weighted average of expert opinions (DeGroot and Mortera 1991; Genest and Zidek 1986; Jose, Grushka-Cockayne and Lichtendahl 2013; Larrick and Soll 2006; McConway 1981; Morris 1977; Stone 1961). In line with that literature, we aggregate beliefs of experts. The key difference is that the literature on opinion pooling does not assume any ordering on the set of outcomes. With the notions of good and bad outcomes thus undefined, there cannot be anything such as a cautious aggregation of beliefs as is key in our contribution.

Our paper fits into the literature that discusses decision-making under imprecise information. In line with other papers (Basili and Chateauneuf 2020; Berger, Emmerling and Tavoni 2016; Berger et al. 2021; Cerreia-Vioglio et al. 2020; Crès, Gilboa and Vieille 2011; Gajdos and Vergnaud 2013; Hansen and Sargent 2001; Heal and Millner 2018; Millner, Dietz and Heal 2013), we borrow from the decision-theoretic literature on ambiguity aversion to model disagreement aversion.[4] The key difference is that to capture aversion to the lack of distribution-consensus, we depart from the standard unanimity principle and follow the dual approach suggested by Bommier (2017).

[4]See Machina and Siniscalchi (2014) for an excellent review of the ambiguity aversion literature.

Our paper is also connected to the general discussion in economics and social choice on how to extend the well-known preference aggregation result in Harsanyi (1955) from risk to uncertainty (Alon and Gayer 2016; Brandl 2020; Danan et al. 2016; Hylland and Zeckhauser 1979; Mongin 1995), and in particular to the debate on how to weaken the Pareto condition (Gilboa, Samet and Schmeidler 2004; Gilboa, Samuelson and Schmeidler 2014; Mongin 2016). Gilboa, Samet and Schmeidler (2004) consider heterogeneity in tastes and beliefs and suggest the Pareto condition be restricted to choices that involve identical beliefs. In our paper, like Stanca (2021), we exclusively focus on heterogeneity in beliefs. All experts are assumed to adopt the risk preferences of the decision-maker, which leaves no room for heterogeneity in tastes. We suggest the Pareto condition be relaxed while keeping a weaker form of unanimity principle based on a property of monotonicity with respect to first-order stochastic dominance. Namely, if all experts believe that a choice α dominates a choice β in the sense of first-order stochastic dominance when looking at the distribution of outcomes, then the decision-maker should prefer α to β.

Finally, the distribution-averaging models we propose constitute a new way of taking cautious decisions when robust scientific knowledge is lacking. As such we contribute to the economic literature on the precautionary principle (Barrieu and Sinclair-Desgagné 2006; Gollier and Treich 2003; Gollier, Jullien and Treich 2000). We offer a convenient and tractable framework where the degree of cautiousness is reflected in a simple averaging function. The framework is flexible and benefits from previous contributions as a number of parametric forms for the averaging function can be directly imported from the literature on ambiguity aversion.

We proceed as follows. Section 1 defines the notion of disagreement aversion, introduces distribution-averaging decision-rules, and clarifies the relation to common expected-utility-averaging models that rest on the unanimity principle. Section 2 compares disagreement aversion and ambiguity aversion in the context of concrete applications and can be read independently. Section 3 concludes.

1 Theoretical framework

1.1 Setting

We consider a decision-maker who has to take decisions in a setting of uncertainty. Decisions will be driven by preferences over uncertain *prospects*. Experts provide scientific knowledge. They share the same objective as the decision-maker but provide conflicting views on the risks at play (i.e., beliefs are heterogeneous). The decision-maker's preferences will naturally depend on the beliefs of all experts, which we call *expertise*. Our analysis bears on *decision-rules* which describe how the expertise shapes the decision-maker's preferences. In the following we provide formal definitions of these concepts.

Prospects Let $X = [X^-, X^+]$, a closed interval of $\mathbb{R}$ with $X^- < X^+$, be the space of *outcomes*. Given a set of *states of the world* Ω, which contains at least three different states, we define (simple) *prospects* as mappings from states of the world to outcomes.[5] Prospects are assumed to have finite images, and for any prospect α, we will denote by K_α the number of values taken by α, and by $(\alpha_1, \ldots, \alpha_{K_\alpha})$ the values taken by α in increasing ordering (i.e. $\alpha_1 < \ldots < \alpha_{K_\alpha}$). For any outcome x, the *sure prospect* with outcome x is the prospect that equals x in all states of the world. With a common abuse of notation, such a sure prospect will also be denoted by x. To avoid confusion, Greek letters α, β will be used for possibly non-sure prospects, while Latin letters such as x will be reserved for outcomes and the corresponding sure prospects.

Expertise We consider N *experts*. Expert $i \in \{1, \cdots, N\}$ holds *belief* P_i, a subjective probability measure on the set of the states of the world Ω. We call *expertise* the list of experts' beliefs $\mathscr{P} = (P_1, \ldots, P_N)$. Making again a slight abuse of notation, the expertise $(P, \ldots, P)$ where all experts share the same belief will also be denoted P. Confusion will be avoided by using calligraphic symbols (like $\mathscr{P}$) for expertise

[5] Prospects are therefore just acts *à la* Savage (1954). We nevertheless opted for a different terminology to avoid the confusion with the notion of acts in the two-stage setting of Anscombe and Aumann (1963) where outcomes are lotteries, i.e. probabilistic objects.

where experts may disagree, and Roman symbols (like P or P_i) for expertise where all experts agree. Given some belief P_i and an outcome x, we denote by $D_\alpha^{P_i}(x)$ the probability that α yields an outcome larger or equal to x, i.e.

$$D_\alpha^{P_i}(x) = P_i\left(\{\omega \in \Omega \,|\, \alpha(\omega) \geq x\}\right).$$

The mapping $x \rightarrow D_\alpha^{P_i}(x)$ is thus the decumulative distribution function of the lottery generated by the prospect α when holding belief P_i.

Decision-rule A decision-rule $\succcurlyeq : \mathscr{P} \mapsto \succcurlyeq^{\mathscr{P}}$ is a mapping that associates an expertise $\mathscr{P}$ with a preference relation (i.e. a weak order) $\succcurlyeq^{\mathscr{P}}$ over the set of prospects. We denote by $\succ^{\mathscr{P}}$ and $\sim^{\mathscr{P}}$ the strict order and indifference relation corresponding to the weak order $\succcurlyeq^{\mathscr{P}}$.

With our notation, the preference relation $\succcurlyeq^{P_i}$ is, formally speaking, the preference relation that the decision-maker uses when all experts have the same beliefs as expert i.

Consensual prospects Central to our analysis is the notion of disagreement aversion (see Section 1.2 below). A notion of disagreement aversion must ceteris paribus rely on a notion of disagreement, or lack of consensus, which could be formalized in different ways.

A first possibility to define consensus involves comparing how experts evaluate prospects in terms of certainty-equivalents. We will say that a prospect α is *utility-consensual* if the prospect's appeal does not depend on the expert the decision-maker might rely on, i.e. when there exists $x \in X$ such that $\alpha \sim^{P_i} x$ for all i. Note that the formulation makes use of the indifference relations $\sim^{P_i}$, which means that the notion of utility-consensual actually depends on the decision-rule considered and in particular on how the decision-maker ranks prospects when experts have identical beliefs. Crucial for our analysis, a prospect may be utility-consensual while experts still fundamentally disagree on the risk implied by such a prospect. An illustration is given in our introductory example: No matter which expert the decision-maker might rely on, the project is seen as good as doing nothing, but one expert thinks that the project is risky and the other not.

As stressed in the introduction, the decision-maker may want to treat cases of complete expert agreement and cases of "spurious agreement" differently. This leads us to introduce a second notion of consensus. We say that a prospect α is *distribution-consensual* when experts agree on the distribution of outcomes it entails (i.e. $D_\alpha^{P_i} = D_\alpha^{P_j}$ for all experts i and j). The concept of disagreement aversion introduced in Section 1.2 will directly rely on the notion of distribution-consensus. In contrast to utility-consensus, distribution-consensus is defined independently of the decision-rule. Note also that all sure prospects are both utility- and distribution-consensual, reflecting that experts' beliefs are irrelevant for the evaluation of sure prospects.

Averagers The specifications we will introduce later on make use of "averager" functionals. Averaging is often conflated with the arithmetic mean. In this paper however, we understand *averaging* in a broad sense. Formally, an *averager* is just a continuous, component-wise strictly increasing[6] function $I : [0,1]^N \to [0,1]$ fulfilling $I(q,\ldots,q) = q$.

Linear averagers (or *weighted means*) take the form $I(q_1,\ldots,q_N) = \sum_{i=1}^N \lambda_i q_i$, with numbers (so-called "weights") $\lambda_1,\ldots,\lambda_N \geq 0$ summing to 1. Linear averagers take a prominent place in the literature on belief aggregation, as they correspond to the ambiguity neutral case. The literature on ambiguity aversion, however, suggests many non-linear averagers. For example, inspired by Klibanoff, Marinacci and Mukerji (2005), one can consider smooth "KMM" averagers of the form:[7]

$$I_S(q_1,\ldots,q_N) = \psi^{-1}\left(\sum_{i=1}^N \lambda_i \psi(q_i)\right) \tag{1}$$

for some increasing smooth function $\psi : [0,1] \to [0,1]$ and some weights $(\lambda_i)_{1\leq i\leq N}$.

[6] By "component-wise" order we mean that for any $\vec{q} = (q_1,\ldots,q_N)$ and $\vec{q}' = (q_1',\ldots,q_N')$ in $[0,1]^N$, the statement $\vec{q} \geq \vec{q}'$ is to be understood as $q_j \geq q_j'$ for all $j \in \{1,\ldots,N\}$. The strict inequality $\vec{q} > \vec{q}'$ is used to mean that $\vec{q} \geq \vec{q}'$ and $\vec{q} \neq \vec{q}'$.

[7] More precisely, such averagers are used in the Second-Order EU model (see e.g., Nau 2006), the most common specification of the model introduced by Klibanoff, Marinacci and Mukerji (2005).

Alternatively, in the spirit of Gilboa and Schmeidler (1989), one can use averagers of the max-min kind,

$$I_M(q_1,\ldots,q_N) = \min_{(\lambda_i)_{1\le i\le N}\in\chi} \sum_{i=1}^{N} \lambda_i q_i, \tag{2}$$

where χ is a closed and convex set of weights. The KMM and max-min are just two specific forms of averagers, but the literature on ambiguity aversion offers many averagers, such as those related to the variational model of Maccheroni, Marinacci and Rustichini (2006) or the model of Siniscalchi (2009).

Risk preferences Throughout the paper, we assume that distribution-consensual prospects are evaluated through a rank-dependent expected utility (RDU) model. RDU is a well-known generalization of the expected utility (EU) theory, introduced by Quiggin (1982).

To gain intuition, recall EU, assuming a utility index $u : X \to [0,1]$. Take a prospect α which is distribution-consensual under the expertise $\mathscr{P} = (P_1,\ldots,P_N)$. Denote $p_k = D_\alpha^{P_1}(\alpha_k)$ for all $1 \le k \le K_\alpha$. Since α is distribution-consensual, one also has $p_k = D_\alpha^{P_i}(\alpha_k)$ for all $1 \le i \le N$ and the expected utility of α under any of the experts' beliefs is given by:

$$EU(\alpha) = \sum_{k=1}^{K_\alpha-1} u(\alpha_k)(p_k - p_{k+1}) + u(\alpha_{K_\alpha})\, p_{K_\alpha}. \tag{3}$$

As emphasized by Castagnoli and LiCalzi (1996), we can present EU in a different (but equivalent) way using a summation by part. Denoting $\sigma_k = u(\alpha_k) - u(\alpha_{k-1}) \ge 0$ for $k \in \{2,\ldots,K_\alpha\}$ and $\sigma_1 = u(\alpha_1)$, equation (3) rewrites:

$$EU(\alpha) = \sum_{k=1}^{K_\alpha} \sigma_k p_k. \tag{4}$$

In order to introduce disagreement aversion, it will prove more convenient to build on formulation (4) than on (3). Moreover, once focusing on (4), it is almost costless to consider the extension to the RDU framework which is obtained by distorting decumulative probabilities p_k through some increasing function f. This leads to the following definition:

Definition 1 (RDU on distribution-consensual prospects)**.** *A decision-rule* $\succcurlyeq$ *is* RDU on distribution-consensual prospects *if there exist two increasing bijections* $f:[0,1]\to[0,1], u:X\to[0,1]$ *such that for any expertise* $\mathscr{P}$*, the weak order* $\succcurlyeq^{\mathscr{P}}$ *restricted on distribution-consensual prospects is represented by:*

$$U^{\mathscr{P}}(\alpha)=\sum_{k=1}^{K_\alpha}\sigma_k f(p_k) \tag{5}$$

where $\sigma_k=u(\alpha_k)-u(\alpha_{k-1})\geq 0$ *for* $k\in\{2,\ldots,K_\alpha\}$*,* $\sigma_1=u(\alpha_1)$ *and* $p_k=D_\alpha^{P_1}(\alpha_k)$*.*

For the sake of conciseness, we will refer to $U^{\mathscr{P}}(\alpha)$ as an "expected utility", even though (5) uses transformed probabilities ($f(p_k)$ instead of p_k). It directly follows from (5) that for any sure prospect x and any expertise $\mathscr{P}$, one has $U^{\mathscr{P}}(x)=u(x)$.[8] Note also that we normalize the utility function by imposing that $\mathrm{Im}(u)=\mathrm{Im}(f)=[0,1]$. This implies that the representation is then unique, but of course without loss of generality.[9]

We could of course have restricted our setting to EU risk preferences (one just needs replace the function f by the identity function in all mathematical expressions that follow). However, since it does not make the analysis, nor the axiomatic construction, significantly more complex, we thought it interesting to keep the additional flexibility of the RDU model which allows extra sensitivity to rare and extreme events to be expressed. This may be valuable for modeling social preferences in the context of uncertainty, especially when considering unlikely but dramatic outcomes (e.g., a major nuclear catastrophe). It is actually interesting to see that the conclusions we will derive regarding the impact of disagreement aversion hold

[8] When α is a sure prospect providing outcome x in all states, one has $K_\alpha=1$, $\alpha_1=x$, and $p_1=1$.

[9] This normalization will also be convenient for comparing expected-utility-averaging and distribution-averaging decision-rules in Section 1.5 since having utility levels and transformed probabilities covering the same interval $[0,1]$ will make it possible to consider identical averager functions, be it to aggregate utilities or to aggregate probabilities.

even when we allow for some extra sensitivity to rare events featured in a RDU framework.

1.2 Disagreement aversion

Disagreement aversion, which is a central concept in our paper, can be seen as a "dislike" for prospects which are not distribution-consensual. It is formally defined as follows:

Definition 2 (Disagreement aversion)**.** *A decision-rule $\succcurlyeq$ is* disagreement-averse *if for every expertise $\mathscr{P}$, every prospect α that is not distribution-consensual, and every sure prospect x, the following implication holds:*

$$\left(x \succcurlyeq^{P_i} \alpha,\, \forall i\right) \Rightarrow x \succ^{\mathscr{P}} \alpha.$$

Note that for decision-rules which are RDU on distribution-consensual prospects, it directly follows from Definition 1 that for all distribution-consensual prospects α, one has $\left(x \succcurlyeq^{P_i} \alpha,\, \forall i\right) \Rightarrow x \succcurlyeq^{\mathscr{P}} \alpha$. Disagreement aversion states that this implication extends to all prospects: Intuitively, if all experts think that α (consensual or not) is not better than x, disagreement among experts cannot make the decision-maker strictly prefer α to x. In addition, disagreement aversion requires that when α is not distribution-consensual, then the decision-maker strictly prefers the sure prospect x.

In the most standard way, we can follow Yaari (1969)'s general approach — initially introduced to compare risk aversion — to compare disagreement aversion. Formally:

Definition 3 (Comparative disagreement aversion)**.** *A decision-rule $\succcurlyeq_A$ exhibits* greater disagreement aversion *than a decision-rule $\succcurlyeq_B$ if for every expertise $\mathscr{P}$, every prospect α and every distribution-consensual prospect β,*

$$\alpha \succ_A^{\mathscr{P}} \beta \Rightarrow \alpha \succ_B^{\mathscr{P}} \beta; \quad \alpha \succcurlyeq_A^{\mathscr{P}} \beta \Rightarrow \alpha \succcurlyeq_B^{\mathscr{P}} \beta$$

and

$$\alpha \sim_A^{\mathscr{P}} \beta \Rightarrow \alpha \succ_B^{\mathscr{P}} \beta \text{ if } \alpha \text{ is not distribution-consensual.}$$

Intuitively, what this definition says is that if according to the decision-rule $\succcurlyeq_A$ the prospect α looks preferable to the distribution-consensual β (despite potential disagreement about the riskiness of α), it must also be the case when considering a decision-rule exhibiting lower disagreement aversion. Moreover, in case α is not consensual, and seen just as good as β when using the more disagreement-averse decision-rule, it must be seen as strictly better than β under the less disagreement averse decision-rule. Note that if $\succcurlyeq_A$ and $\succcurlyeq_B$ are comparable in terms of disagreement aversion, in the sense that one exhibits more disagreement aversion than the other, they must necessarily agree on the ranking of distribution-consensual prospects (see proof in Appendix A.3). This parallels the literature on risk aversion, where agents are comparable in terms of risk aversion only if they agree on the ranking of deterministic outcomes (see Kihlstrom and Mirman 1974).

Violation of the Pareto condition Disagreement aversion directly conflicts with the Pareto (or unanimity) condition. The fundamental reason is that utility-consensual prospects are not necessarily distribution-consensual. In order to formally stress the tension between disagreement aversion and the Pareto condition, we introduce the notion of Paretian decision-rule.

Definition 4 (Paretian decision-rule)**.** *A decision-rule $\succcurlyeq$ is* Paretian *if for every expertise $\mathscr{P} = (P_1, \ldots, P_N)$ and all prospects α, β:*

$$\left(\beta \succcurlyeq^{P_i} \alpha, \forall i\right) \Rightarrow \beta \succcurlyeq^{\mathscr{P}} \alpha.$$

The following result immediately follows.

Lemma 1. *There is no decision-rule which is Paretian, RDU on distribution-consensual prospects, and exhibits disagreement aversion.*

Proof. Suppose that the decision-rule is RDU on distribution-consensual prospects. One can construct an expertise $\mathscr{P} = (P_1, \ldots, P_N)$, a non distribution-consensual prospect α and a sure prospect x such that $x \sim^{P_i} \alpha$, $\forall i$, e.g. using the beginnings of the proofs of Proposition 1 (in Appendix A.1) and Proposition OA.1 (in online Appendix OA.1). If the decision rule is Paretian we then have $x \sim^{\mathscr{P}} \alpha$, contradicting $x \succ^{\mathscr{P}} \alpha$, and thus disagreement aversion. ■

Lemma 1 may naturally lead to questioning whether the most appealing property is disagreement aversion or the Pareto condition. If the decision-maker could be sure that one expert is right, then the Pareto condition would seem natural. However, if the decision-maker interprets the divergence in experts' beliefs as reflecting some fundamental imprecision in knowledge (meaning that most likely none of the experts is right), then disagreement aversion seems a natural way to model cautious decision-making. Our introductory example can help the reader make up their own mind about these respective properties. Indeed, (strict) preference for the no-project option would be required under disagreement aversion, and ruled out under the Pareto condition.

Lemma 1 makes it clear that in order to model disagreement aversion, we need to relax the Pareto condition. An indirect consequence is that we cannot build upon the major ambiguity models discussed in the literature (e.g., in Machina and Siniscalchi 2014 and Strzalecki 2013), as they rely on a monotonicity axiom which in our setting translates to the Pareto condition. An exception in the ambiguity aversion literature is Bommier (2017), who explicitly relaxes the monotonicity axiom. We will borrow from this paper to suggest a class of decision-rules exhibiting disagreement aversion.

1.3 Distribution-averaging decision-rules

The following definition introduces the class of *distribution-averaging decision rules* for which disagreement may matter, possibly to exhibit disagreement-aversion, disagreement-loving or some ambiguous patterns of disagreement sensitivity. Below the definition, we explain how such a specification may be derived from a formal set of axioms.

Definition 5 (Distribution-averaging decision-rule)**.** *A decision-rule* $\succcurlyeq$ *is* distribution-averaging *with representation* (u, f, I) *if there exist two increasing bijections* $u : X \to [0,1]$*,* $f : [0,1] \to [0,1]$*, and an averager* $I : [0,1]^N \to [0,1]$ *such* *that for any expertise $\mathscr{P}$, the weak order $\succcurlyeq^{\mathscr{P}}$ is represented by:*

$$U^{\mathscr{P}}(\alpha) = \sum_{k=1}^{K_\alpha} \sigma_k I\left(f\left(p_k^1\right), \ldots, f\left(p_k^N\right)\right) \tag{6}$$

where $\sigma_k = u(\alpha_k) - u(\alpha_{k-1}) \geq 0$ for $k \in \{2, \ldots, K_\alpha\}$, $\sigma_1 = u(\alpha_1)$ and $p_k^i = D_\alpha^{P_i}(\alpha_k)$.

We will refer to u as the utility index and to f as the probability transformation function.

Remark 1. Because the range of u is normalized to $[0,1]$, the representation is unique, in the sense that a decision-rule admits only one (u, f, I) representation. A distribution-averaging decision rule is necessarily RDU on consensual prospect, with representation as in equation (5).

Remark 2. Defining $\widetilde{I} : \left(p^1, \ldots, p^N\right) \mapsto \widetilde{I}\left(p^1, \ldots, p^N\right) = f^{-1}\left(I\left(f\left(p^1\right), \ldots, f\left(p^N\right)\right)\right)$, equation (6) simply rewrites:

$$U^{\mathscr{P}}(\alpha) = \sum_{k=1}^{K_\alpha} \sigma_k f\left(\widetilde{I}\left(p_k^1, \ldots, p_k^N\right)\right). \tag{7}$$

Thus, as shown in Figure 1, a distribution-averaging decision-rule can be seen as one where the values of the decumulative distribution functions provided by each expert are "averaged" through the function $\widetilde{I}$ in order to provide some equivalent decumulative distribution function (i.e. some equivalent belief). The prospect is then evaluated with the RDU specification using this equivalent decumulative distribution function. It is worth noting that the way some belief P_i ends up giving a decumulative distribution function depends on the prospect that is evaluated (mathematically, p_k^i depends on both P_i and α). Thus, aggregating decumulative distribution functions associated with a given prospect is not equivalent to a mere aggregation of beliefs, which would be made independently of whether states provide high or low pay-offs.[10]

[10]This is a fundamental difference with the linear opinion pooling where the decision-maker averages linearly expert beliefs to form their own belief, independently of the prospect to be evaluated.

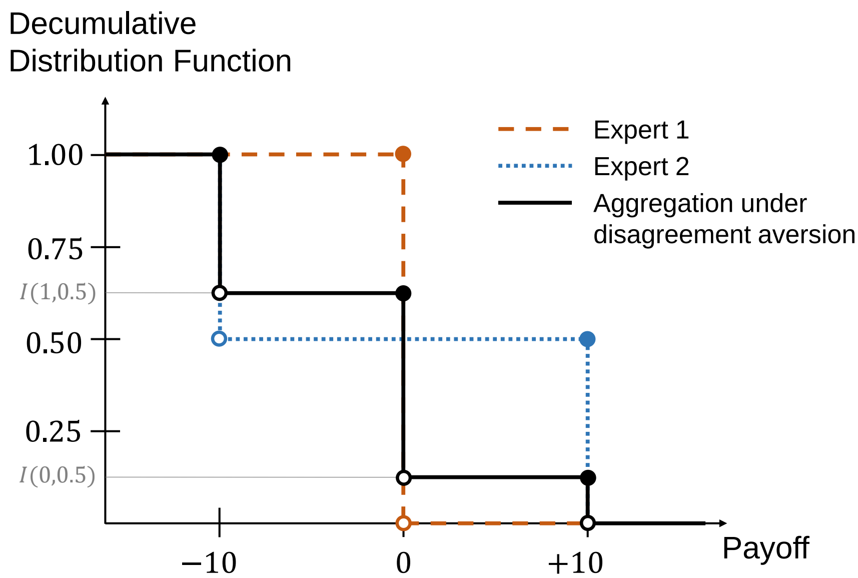

Figure 1 – Aggregation of experts' beliefs in the introductory example.

Axiomatization An axiomatization of distribution-averaging decision-rules is provided in online Appendix OA.2. The main axioms are:

- Monotonicity with respect to first-order stochastic dominance (M-FSD). This axiom is weaker than the monotonicity axiom assumed in models based on the unanimity principle (or Pareto condition), which rules out disagreement aversion. Our M-FSD axiom only requires that a decision-rule prefers a prospect α to a prospect β ($\alpha \succcurlyeq \beta$) if α dominates β with respect to first-order stochastic dominance, i.e. $D_{\alpha}^{P_i} \geq D_{\beta}^{P_i}, \forall i$, while the Pareto condition requires that it be the case as soon as $U^{P_i}(\alpha) \geq U^{P_i}(\beta), \forall i$.

- RDU on distribution-consensual prospects. Equivalently, this axiom can be replaced by comonotonic mixture independence (see Chateauneuf 1999).

- The comonotonic sure-thing principle. This specifies that a common outcome of two prospects at some state of the world can be changed without impacting the comparison between those prospects, as long as the change does not affect the rankings of each prospect's outcomes.

- A usual continuity axiom.

- An axiom of "level independence", which restricts the type of disagreement aversion that the decision-rule can exhibit; namely, we require that for prospects with only two possible outcomes, the way the disagreement between experts is resolved does not depend on the outcomes but only on the probabilities. This axiom is introduced to obtain a simple representation. One could however relax it, to obtain a broader class of distribution-averaging decision-rules.[11]

1.4 Characterization of disagreement aversion

We first state what is required to have a decision-rule that exhibits disagreement aversion.

Proposition 1. *A distribution-averaging decision-rule with representation (u, f, I) exhibits disagreement aversion if and only if the averager I is such that for any matrix $(q_k^i)_{1\leq k\leq K, 1\leq i\leq N} \in [0,1]^{N\times K}$ such that $k \geq l \Rightarrow q_k^i \leq q_l^i$, and any vector $(\sigma_1, \ldots, \sigma_K) \in [0,1]^K$ such that $\sum_{k=1}^K \sigma_k \leq 1$, one has:*

$$\sum_{k=1}^{K} \sigma_k I\left(q_k^1, \ldots, q_k^N\right) \leq \max_{1\leq i\leq N} \sum_{k=1}^{K} \sigma_k q_k^i$$

where the inequality is strict whenever one has $q_k^i \neq q_k^j$ and $\sigma_k > 0$ for some i, j, k.

Proof. See Appendix A.1. ■

To our knowledge, there is no available name for the property stated in the above proposition. One can, however, provide simpler conditions on I that are sufficient (but not necessary) to obtain disagreement aversion.

Proposition 2. *Each of the conditions below is sufficient to imply disagreement aversion:*

- **Condition 1**: *There exist numbers $(\lambda_i)_{1\leq i\leq N}$ (also called weights), which are non-negative and sum to 1 such that $I(p_1, \ldots, p_N) \leq \sum_{i=1}^N \lambda_i p_i$ for all $(p_1, \ldots, p_N) \in [0,1]^N$, with a strict inequality when $p_i \neq p_j$ for some i and j.*

[11] One would then obtain a representation akin to the one given in Bommier (2017), Theorem 2.

- **Condition 2**: *The function* $(p_1,\ldots,p_N) \to I(p_1,\ldots,p_N)$ *is strictly concave except on the diagonal.*[12]

Proof. See Appendix A.2. ∎

Condition 1 is convenient, as it makes it possible to borrow from the literature on ambiguity aversion where such an inequality (considered then for averagers that average utilities rather than probabilities) is the one that usually defines ambiguity aversion (e.g., Cerreia-Vioglio et al. 2011).

The last result of this section is about comparative disagreement aversion.

Proposition 3. *Consider two distribution-averaging decision-rules* $\succcurlyeq_A$ *and* $\succcurlyeq_B$ *with representation* (u_A, f_A, I_A) *and* (u_B, f_B, I_B)*. Then* $\succcurlyeq_A$ *exhibits greater disagreement aversion than* $\succcurlyeq_B$ *if and only if* $u_A = u_B$*,* $f_A = f_B$ *and* $I_A(\vec{p}) < I_B(\vec{p})$ *for every non-constant vector* $\vec{p} = (p_1,\ldots,p_N) \in [0,1]^N$.

Proof. See Appendix A.3. ∎

The degree of disagreement aversion is thus fully determined by the averager, with "smaller" averagers yielding more disagreement-averse decision-rules. This shows similarity with results from the ambiguity aversion literature, where ambiguity aversion is driven by how "pessimistic" the averagers (which aggregate utilities in that literature) are. This is convenient, as we can readily import knowledge from the ambiguity literature to determine what it means to increase disagreement aversion. For example, if we consider smooth KMM averagers as in equation (1), we readily know that increasing disagreement aversion is equivalent to increasing the concavity of the function ψ. Similarly, if we consider max-min averagers defined as in equation (2), we know that increasing disagreement aversion involves using a larger set of weights χ.

1.5 Relation with Paretian decision-rules

In order to explain how our work relates to previous contributions, we introduce a class of decision-rules that rest on the unanimity principle.

[12] By strictly concave except on the diagonal, we mean that the strict concavity inequality holds for any pair of vectors of which at least one is not constant.

Definition 6 (Expected-utility-averaging decision-rule). *A decision-rule* $\succcurlyeq$ *is* expected-utility-averaging *with representation* (u, f, I) *if there exist two increasing bijections* $u : X \to [0,1]$, $f : [0,1] \to [0,1]$, *and an averager* $I : [0,1]^N \to [0,1]$ *such that for any expertise* $\mathscr{P}$, *the weak order* $\succcurlyeq^{\mathscr{P}}$ *is represented by:*

$$U^{\mathscr{P}}(\alpha) = I\left(\sum_{k=1}^{K_\alpha} \sigma_k f\left(p_k^1\right), \ldots, \sum_{k=1}^{K_\alpha} \sigma_k f\left(p_k^N\right)\right) \tag{8}$$

where $\sigma_k = u(\alpha_k) - u(\alpha_{k-1}) \geq 0$ *for* $k \in \{2, \ldots, K_\alpha\}$, $\sigma_1 = u(\alpha_1)$ *and* $p_k^i = D_\alpha^{P_i}(\alpha_k)$.

The qualification of *expected-utility-averaging* comes from the fact that (8) can be rewritten as $U^{\mathscr{P}}(\alpha) = I\left(U^{P_1}(\alpha), \ldots, U^{P_N}(\alpha)\right)$ where $U^{P_i}(\alpha) = \sum_{k=1}^{K_\alpha} \sigma_k f\left(p_k^i\right)$ is the RDU utility obtained when holding belief P_i. It directly follows from the monotonicity of the averager I that such decision-rules are Paretian in the sense of Definition 4. This class of expected-utility-averaging models embeds all major models listed in Machina and Siniscalchi (2014). The representation (8) is actually very similar to that of Cerreia-Vioglio et al. (2011), a very general specification that encompasses most ambiguity models.[13] In those ambiguity models, the averager I is what characterizes ambiguity aversion.

Comparing equation (6) and equation (8), we see that distribution-averaging and expected-utility-averaging decision-rules only differ by the stage where the averager is applied. In the former case, I aggregates the distributions and expected utility is computed afterwards. In the latter, one starts by computing expected utility levels, which are then aggregated through the averager I. Aggregating distributions first (as with distribution-averaging decision-rules) makes it possible to exhibit aversion for divergences in probabilities reported by experts and thus to exhibit disagreement aversion. On the other hand, computing expected utilities first (as with expected-utility-averaging decision-rules) may lead to losing track of some divergence in expert opinions and thus prevents the expression of disagreement aversion.

When the averager I is linear, both approaches are equivalent, and the decision-rule is both distribution-averaging and expected-utility-averaging. The reciprocal is

[13]Compared to the Monotonic, Bernoullian and Archimedean (MBA) model of Cerreia-Vioglio et al. (2011), our expected-utility-averaging model evaluates risk using RDU instead of EU. In so doing, we drop the Bernoullian assumption in MBA and consider an even broader class of models.

also true:

Proposition 4. *A decision-rule is both distribution-averaging and expected-utility-averaging if and only if it admits a (distribution-averaging or expected-utility-averaging) representation (u, f, I) with a linear averager I.*

This result, whose proof is given in Appendix A.4, shows that being simultaneously distribution-averaging and expected-utility-averaging requires having a linear averager I. This amounts to linear pooling and is usually considered as characterizing ambiguity neutrality. The distribution-averaging and expected-utility-averaging approaches differ as soon as non-linear averagers are assumed. Interestingly, they can be compared in terms of ambiguity aversion, a fundamental concept in the ambiguity literature. In our setting, comparative ambiguity aversion can be defined as follows:

Definition 7 (Comparative ambiguity aversion)**.** *A decision-rule $\succcurlyeq_A$ exhibits* greater ambiguity aversion *than a decision-rule $\succcurlyeq_B$ if for every expertise $\mathscr{P}$, every prospect α and every distribution-consensual prospect β,*

$$\alpha \succ_A^{\mathscr{P}} \beta \Rightarrow \alpha \succ_B^{\mathscr{P}} \beta; \quad \alpha \succcurlyeq_A^{\mathscr{P}} \beta \Rightarrow \alpha \succcurlyeq_B^{\mathscr{P}} \beta$$

and

$$\alpha \sim_A^{\mathscr{P}} \beta \Rightarrow \alpha \succ_B^{\mathscr{P}} \beta \text{ if } \alpha \text{ is not utility-consensual.}$$

This definition is very similar to Definition 3 about disagreement aversion but relies on a different notion of consensus. While Definition 3 refers to the aversion for the lack of distribution consensus, Definition 7 refers to the aversion for a lack of utility consensus.[14]

Remark 3. Ghirardato and Marinacci (2002) define comparative ambiguity aversion in a weak sense (such that any rule is more ambiguity averse than itself). By adding a third condition (about utility-consensual prospects), our definition of comparative ambiguity aversion is strong, and characterized by strict (rather than non-strict)

[14]The first part of Definition 7 still refers to the notion of distribution-consensus, so that the decisions rules $\succcurlyeq_A$ and $\succcurlyeq_B$ necessarily agree on the ranking of distribution-consensual prospects. The notion of utility-consensus used in the second part of the definition which, in principle, could depend on the decision rule that is considered, is then the same for $\succcurlyeq_A$ and $\succcurlyeq_B$.

inequalities. We add the third condition to make the parallel with comparative disagreement aversion (Definition 3).

We first state a result (already well-known for expected-utility-averaging models – see Ghirardato, Maccheroni and Marinacci 2004) showing that for both distribution-averaging decision-rules and expected-utility-averaging decision-rules, the level of ambiguity aversion is dictated by the averager I. Formally:

Proposition 5. *Consider two decision-rules $\succcurlyeq_A$ and $\succcurlyeq_B$, which are either both distribution-averaging or both expected-utility-averaging with representations (u_A, f_A, I_A) and (u_B, f_B, I_B). Then $\succcurlyeq_A$ exhibits greater ambiguity aversion than $\succcurlyeq_B$ if and only if $u_A = u_B$, $f_A = f_B$ and $I_A(\vec{p}) < I_B(\vec{p})$ for every non-constant vector $\vec{p} = (p_1, \ldots, p_N)$.*

Proof. See Appendix A.3. ■

Proposition 5 shows how two decision-rules of the same kind can be compared in terms of ambiguity aversion.

The following result highlights that it is also possible to compare distribution-averaging decision-rules with expected-utility-averaging decision-rules. Namely, under a concavity condition on the averager I, which is usually assumed in applications, we show that a distribution-averaging decision-rule exhibits greater ambiguity aversion and greater disagreement aversion than the expected-utility-averaging decision-rule, which uses the same utility index, the same probability transformation function and the same averager. More precisely, the result is valid only when we exclude the special case of prospects whose outcomes are all extremal (i.e. $\alpha_k \in \{X^-, X^+\}, \forall k$), because both rules coincide on these prospects. For the purpose of the proposition, we thus say that a decision-rule exhibits greater disagreement aversion (resp. ambiguity aversion) *except on prospects with only extremal outcomes* if it fulfills a slightly weaker version of Definition 3 (resp. Definition 7) where α is replaced by any prospect with at least one non-extremal outcome (i.e. $0 < \sigma_k < 1$ for some k).

Proposition 6. *Consider a distribution-averaging decision-rule $\succcurlyeq_{DA}$ with representation (u, f, I) and the expected-utility-averaging decision rule $\succcurlyeq_{UA}$ with representation (u, f, I) for the same functions u, f, and I. If I is strictly concave except on the* *diagonal then $\succcurlyeq_{DA}$ exhibits greater ambiguity aversion and greater disagreement aversion than $\succcurlyeq_{UA}$ except on prospects with only extremal outcomes.*

Proof. See Appendix A.5. ■

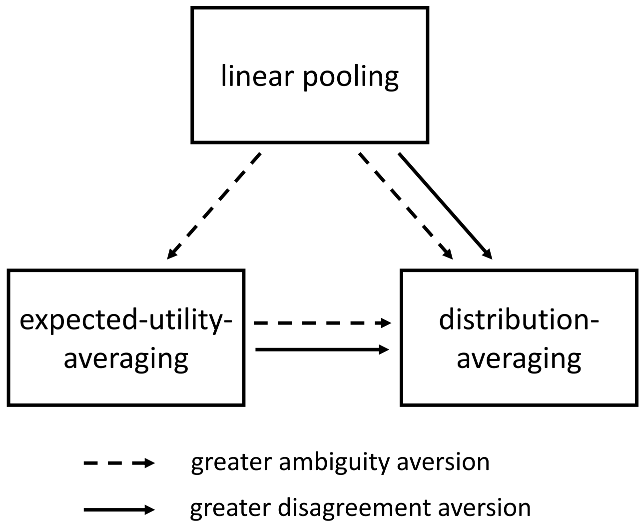

Figure 2 – The relation between decision-rules when I is strictly concave.

Combining several of the above results, Figure 2 provides a global picture on how distribution-averaging and expected-utility-averaging decision-rules relate to each other. Starting from the linear pooling case where the decision-maker's utility is a linear average of experts' utilities, the distribution-averaging and expected-utility-averaging frameworks offer two different ways to introduce ambiguity aversion. The expected-utility-averaging framework does so while preserving the unanimity principle. This prevents however the expression of disagreement aversion. The distribution-averaging framework allows for further ambiguity aversion, by aggregating probabilities first, and then computing expected utilities. This is what produces disagreement aversion.

Our motivation for formalizing the notion of disagreement aversion and suggesting the distribution-averaging model is not purely theoretical but also rooted in the challenge of expert disagreement in concrete applications. Accounting for disagreement aversion may indeed justify more precautionary decisions when experts disagree and provide insights on how scientific expertise should be used for decision-making. We develop these aspects in the following section.

2 Applications

We consider a decision-maker whose ex-post utility $u(a,\omega)$ depends on an action $a \in [\underline{a},\overline{a}]$ and a contingency $\omega \in [\underline{\omega},\overline{\omega}]$ over which she has no control. We assume that the function $(a,\omega) \mapsto u(a,\omega)$ is increasing in ω, which is without loss of generality when contingencies can be ranked from "adverse" to "favorable" independently of the decision-maker's action. Ex-ante, when the decision-maker has to decide about a, the value of the contingency ω is uncertain. There are N experts who provide potentially diverging views about the distribution of ω. To make the link with the theory section, we assume that the distributions provided by the experts have finite support, with $\omega \in \{\omega_1,\dots,\omega_K\}$ where $(\omega_k)_{1\le k\le K}$ is an increasing sequence of real numbers ($k<l \Rightarrow \omega_k<\omega_l$). Each expert $i \in \{1,\dots,N\}$ provides a probability vector $(\pi_k^i)_{1\le k\le K}$, where π_k^i is the probability of $\omega=\omega_k$ according to expert i. We assume that there is some disagreement among experts ($\pi_k^i \neq \pi_k^j$ for some i, j, k).

2.1 A general result on the impact of disagreement aversion

We assume here that decision-makers use distribution-averaging decision-rules as introduced in Definition 5, with the only difference that we no longer require the utility to be normalized to values in $[0,1]$. Since our aim is to discuss the impact of disagreement aversion, we will consider two decision-makers, A and B, who only differ by their averagers I_A and I_B. Formally speaking, the decision-maker $\tau = A,B$ has the following program:

$$\max_{a\in[\underline{a},\overline{a}]} \sum_{k=1}^{K} \sigma_k I_\tau \left(f\left(p_k^1\right),\dots,f\left(p_k^N\right)\right), \tag{9}$$

where $p_k^i = \sum_{l=k}^{K} \pi_l^i$, $\sigma_k = u(a,\omega_k) - u(a,\omega_{k-1})$ for $k>1$ and $\sigma_1 = u(a,\omega_1)$. The function f is an increasing bijection over $[0,1]$ and the averager I_τ is a continuous and componentwise strictly increasing function $I_\tau : [0,1]^N \to [0,1]$ fulfilling

$I_\tau(q,\ldots,q) = q$. We assume that the decision-maker's problem always has an interior solution denoted $a^*_\tau \in (\underline{a}, \overline{a})$.

Proposition 7. *Assume that $(a,\omega) \mapsto u(a,\omega)$ is twice continuously differentiable and strictly concave in a, and that distribution-averaging decision-maker A exhibits greater disagreement aversion than B, i.e. $I_A(p_1,\ldots,p_N) < I_B(p_1,\ldots,p_N)$ for all non-constant vectors $(p_1,\ldots,p_N) \in [0,1]^N$. Then:*

- *If $\frac{\partial^2 u(a,\omega)}{\partial a\, \partial \omega} > 0$ for all $a \in [\underline{a}, \overline{a}]$ and $\omega \in [\underline{\omega}, \overline{\omega}]$, then $a^*_A < a^*_B$.*
- *If $\frac{\partial^2 u(a,\omega)}{\partial a\, \partial \omega} < 0$ for all $a \in [\underline{a}, \overline{a}]$ and $\omega \in [\underline{\omega}, \overline{\omega}]$, then $a^*_A > a^*_B$.*

Proof. See Appendix A.6. ■

This proposition shows that when the cross-derivative of the function u has a constant sign, an increase of disagreement aversion has an unambiguous impact on the optimal action. The reason is as follows. If the cross-derivative is positive, an increase of a widens the distance between any two possible ex-post utility levels. This makes the disagreement between experts (i.e. the difference in their distribution functions) more significant, which, in turn, implies that an increase in disagreement aversion leads to a decrease in the optimal action a. Vice versa for a negative cross-derivative. It is worth noting that to obtain this result we did not have to assume a particular functional form for the averager. It thus holds for KMM types or max-min averagers, as with other averagers that can be found in the ambiguity aversion literature.

Below, we illustrate the general result in Proposition 7 with two examples, climate policy and precautionary savings.

2.1.1 A climate mitigation example

As highlighted by the literature (e.g., Meinshausen et al. 2009), experts in climate physics substantially disagree on climate sensitivity, i.e. by how much global average temperatures increase as a result of increased greenhouse gas levels. Climate sensitivity is a key parameter for determining the optimal level of greenhouse gas abatement. Should disagreement among experts lead to choosing a higher or lower emission abatement target?

In order to give insights on this question we consider a two-period model. In period 1, the decision-maker is endowed with wealth w_1 of which an amount $C(a)$, increasing and convex in a, can be taken to finance an abatement level $a \in [0, \overline{a}]$, leaving $w_1 - C(a)$ for consumption. In period 2, consumption equals wealth w_2 minus climate-related damages. These damages depend on the abatement level a chosen in period 1 and climate sensitivity $\theta \in [\underline{\theta}, \overline{\theta}]$, which is ex-ante uncertain. We denote by $(a, \theta) \mapsto D(a, \theta)$ the damage function and assume that $\frac{\partial D}{\partial a} < 0$, $\frac{\partial D}{\partial \theta} > 0$, $\frac{\partial^2 D}{\partial a^2} > 0$ and $\frac{\partial^2 D}{\partial a \, \partial \theta} < 0$; the last inequality means that abatement is more efficient in reducing damages when climate sensitivity is high. The ex-post utility is given by:

$$u(a, \theta) = v(w_1 - C(a)) + \beta v(w_2 - D(a, \theta))$$

where v is the (increasing and concave) instantaneous utility function and $\beta > 0$ the time discount factor. One can easily check that:

$$\frac{\partial u(a, \theta)}{\partial \theta} < 0 \ ; \ \frac{\partial^2 u(a, \theta)}{\partial a^2} < 0 \ ; \ \frac{\partial^2 u(a, \theta)}{\partial a \, \partial \theta} > 0.$$

To conform with the setting of Proposition 7, we set $\omega = -\theta$, so that utility increases with ω. It then directly follows from Proposition 7 that if the solution is interior, an increase of disagreement aversion leads to a strict increase of emission abatement. The interpretation is simple: Both in terms of ex-post wealth and ex-post utility, the marginal benefit of emission abatement is higher in bad states of nature than in good states of nature. Thus abatement is able to reduce the distance between ex-post utilities, which implies that greater disagreement aversion leads to an increase of optimal abatement.

2.1.2 A precautionary savings example

We now consider a standard precautionary savings problem in a two-period setting. The decision-maker receives income y in period 1 and has uncertain income $\omega \in [\underline{\omega}, \overline{\omega}]$ in period 2. The decision-maker chooses the amount $a \in [0, \overline{a}]$ saved in period 1. We assume a deterministic rate of interest, r, so that saving a in period 1 yields $(1 + r)a$ in period 2. Choosing an amount of saving a and receiving income

ω in period 2 provides an (ex-post) intertemporal utility:

$$u(a,\omega) = v(y-a) + \beta v(\omega + (1+r)a)$$

where v is the (increasing and strictly concave) instantaneous utility function and $\beta > 0$ the time discount factor. One has:

$$\frac{\partial u(a,\omega)}{\partial \omega} > 0 \;\; ; \;\; \frac{\partial^2 u(a,\theta)}{\partial a^2} < 0 \;\; ; \;\; \frac{\partial^2 u(a,\omega)}{\partial a \, \partial \omega} < 0.$$

Thus a direct application of Proposition 7 shows that if the solution is interior, an increase in disagreement aversion leads to a strict increase of precautionary savings. While the marginal benefit of precautionary savings is similar across states of nature in terms of ex-post wealth, it is higher in bad states of nature than in good states of nature in terms of ex-post utility. Thus precautionary savings reduce the distance between ex-post utilities, which implies that an increase of disagreement aversion leads to higher precautionary savings.

2.2 Expected-utility-averaging versus distribution-averaging

This section highlights how our model differs from common models of ambiguity aversion in applications. As explained in Section 1.5, distribution-averaging decision-rules and common models of ambiguity aversion differ as the former suggests aggregating probabilities first, and then computing expected utilities, while the latter suggests computing expected utilities first, and then aggregating utilities. We will show in this section, that both procedures may yield qualitatively similar results in some cases, but very different, possibly contrary results, in others.

2.2.1 The impact of ambiguity aversion under expected-utility-averaging

We consider here two decision-makers, A and B who use expected-utility-averaging decision-rules but only differ by the averagers I_A and I_B they are using. With the setting described at the beginning of Section 2, the decision-maker $\tau = A, B$ has the following program:

$$\max_{a \in [\underline{a}, \overline{a}]} I_\tau \left(\sum_{k=1}^{K} \sigma_k f\left(p_k^1\right), \ldots, \sum_{k=1}^{K} \sigma_k f\left(p_k^N\right) \right), \tag{10}$$

where the σ_k and the p_k^i are defined as in Section 2.1. We assume that the decision-maker's program always admits an interior solution $a_\tau^* \in (\underline{a}, \overline{a})$.

We know from Proposition 5 that decision-maker A exhibits greater ambiguity aversion than decision-maker B if $I_A(u_1, \ldots, u_N) < I_B(u_1, \ldots, u_N)$ for non-constant vectors of utility levels $(u_1, \ldots, u_N)$. This is however generally insufficient to derive general conclusions about the impact on the optimal action. As noticed by Gollier (2011), Millner, Dietz and Heal (2013), Berger, Emmerling and Tavoni (2016), Berger (2014) or Peter (2019), some results can nevertheless be provided under specific assumptions regarding the averagers I_A and I_B and how experts' beliefs compare. In particular:

Proposition 8. *Assume that:*

- *I_A and I_B have the smooth "KMM" form (1) with twice continuously differentiable, increasing and concave functions ψ_A and ψ_B.*
- *ψ_A is more concave than ψ_B (so that expected-utility-averaging decision-maker A is more ambiguity-averse than B).* [15]
- *$p_k^1 = \sum_{l=k}^{K} \pi_l^1 \leq \ldots \leq p_k^N = \sum_{l=k}^{K} \pi_l^N$ for all k (meaning that the experts' beliefs can be ordered in terms of first-order stochastic dominance).*
- *The function $(a, \omega) \mapsto u(a, \omega)$ is twice continuously differentiable, strictly concave in a and such that $\frac{\partial^2 u(a,\omega)}{\partial a\, \partial \omega} > 0$ (resp. $\frac{\partial^2 u(a,\omega)}{\partial a\, \partial \omega} < 0$) for all $a \in [\underline{a}, \overline{a}]$ and $\omega \in [\underline{\omega}, \overline{\omega}]$.*

Then $a_A^* < a_B^*$ (resp. $a_A^* > a_B^*$).

Proof. See Appendix A.7. ■

The result stated in Proposition 8 looks qualitatively similar to the one stated in Proposition 7 regarding disagreement-averse decision-rules. The interpretation

[15] By "more concave", we mean that there exists an increasing and strictly concave function $h : \mathbb{R} \to \mathbb{R}$ such that $\psi_A = h(\psi_B)$.

is also qualitatively similar, with the difference that we now consider the dispersion of ex-ante (expected) utilities and not the dispersion of ex-post utilities. If the cross-derivative of the function u is positive and experts are ordered in the sense of first-order stochastic dominance, an increase in a widens the distance between the expected utility levels provided by any two experts. Thus, the higher the action a the more significant the lack of unanimity between experts, which implies that an increase of ambiguity aversion leads to a decrease of the optimal a. This result requires however an important additional assumption relative to Proposition 7. Indeed Proposition 8 assumes that experts' beliefs are comparable in terms of first-order stochastic dominance. This means that experts can unambiguously be ranked in terms of optimism (or pessimism), ruling out cases where some experts look more optimistic than others on some aspects but less optimistic on other aspects. For example, it rules out the case discussed in our introductory example where expert 2 looks more optimistic than expert 1 as she predicts that the project may generate positive returns but also looks more pessimistic than expert 1 as she foresees cases where the project would generate negative returns. In such a case, Proposition 8 has no bite.

The following section provides an example where expected-utility-averaging and distribution-averaging decision-rules lead to opposite conclusions.

2.2.2 When utility-averaging and distribution-averaging models differ

Let us consider again the two examples described in Sections 2.1.1 and 2.1.2, while specifying further the utility functions and expert beliefs.

For the climate mitigation example, assume that the instantaneous utility function v and the probability transformation function f are the identity function and the discount factor is $\beta = 1$. Consider the case where $w_1 = 2$, $C(a) = 1.5a^2$, $w_2 = 3.5$, $D(a,\theta) = 0.5(1+\theta-a)^2$, and $a \in [0, 0.9]$. Assume that θ can only take one of the following three values $\theta_1 = 1.5$, $\theta_2 = 1$ or $\theta_3 = 0$ (i.e., $\omega_1 = -1.5$, $\omega_2 = -1$ and $\omega_3 = 0$). With all these assumptions, the conditions assumed in the climate mitigation example described in Section 2.1.1 are satisfied.

For the precautionary savings example, assume that the instantaneous utility function is quadratic, with $v(c) = c - \frac{1}{8}c^2$, the discount factor is $\beta = 1$, and the probability transformation function f is the identity function. Consider the case where the first-period income is $y = 2$, while the second-period income can only take one of the following three values $\omega_1 = 0$, $\omega_2 = 1$ or $\omega_3 = 3$. Assume that the rate of interest is $r = 0$ and $a \in [0, 0.9]$. With all these assumptions, the conditions assumed in the precautionary savings example described in Section 2.1.2 are satisfied.

In both examples, assume that there are only two experts. According to expert 1, the likelihood of the three values ω_1, ω_2 and ω_3 are given by the probability vector $(\pi_1^1, \pi_2^1, \pi_3^1) = (0, 1, 0)$. Expert 2 disagrees and holds belief $(\pi_1^2, \pi_2^2, \pi_3^2) = (0.6, 0, 0.4)$. Finally, consider the averager:

$$(x_1, x_2) \in \mathbb{R} \mapsto I(x_1, x_2) = -\frac{1}{\lambda} \log\left(\frac{1}{2}e^{-\lambda x_1} + \frac{1}{2}e^{-\lambda x_2}\right) \tag{11}$$

with $\lambda \geq 0$.[16] This averager is defined over $\mathbb{R}^2$ and will be used either to aggregate probabilities (for distribution-averaging decision-rules) or to aggregate utility values (for expected-utility-averaging decision-rules). In all cases, increasing λ involves increasing ambiguity aversion. This also generates an increase in disagreement aversion when using distribution-averaging decision-rules.

Denote by $a^*_{DA,\lambda}$ the optimal action (i.e., optimal abatement in the climate example, and optimal saving in the second example) when using the distribution-averaging decision-rule with averager (11) and by $a^*_{UA,\lambda}$ the optimal action when using an expected-utility-averaging decision-rule with that same averager. It can easily be verified that the solutions $a^*_{DA,\lambda}$ and $a^*_{UA,\lambda}$ are interior for any λ.

Proposition 9. *In both examples, the optimal action $a^*_{DA,\lambda}$ is increasing with λ (i.e. with disagreement aversion) and the optimal action $a^*_{UA,\lambda}$ is decreasing with λ (i.e. with ambiguity aversion).*

Proof. See Appendix A.8. ∎

The above result, which would actually extend to all averagers of the KMM kind, indicates that we have here examples where disagreement aversion leads to

[16]The limit case when $\lambda \to 0$, provides $I(x_1, x_2) = \frac{1}{2}x_1 + \frac{1}{2}x_2$, corresponding to a linear opinion pooling with symmetric weights.

greater abatement or larger precautionary savings while ambiguity aversion in the expected-utility-averaging models is found to have the opposite result.

Figure 3 shows how $a^*_{DA,\lambda}$ and $a^*_{UA,\lambda}$ vary with λ in each example. We also report the optimal actions a^*_1 and a^*_2 that would be chosen if relying exclusively on the beliefs of expert 1 or on those of expert 2.

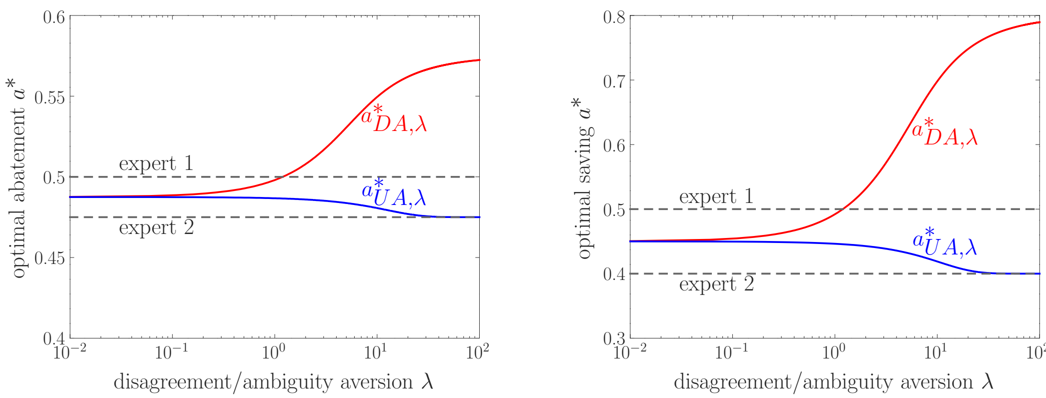

Figure 3 – Optimal abatement (left figure) and optimal saving (right figure) for the distribution-averaging decision-rule ($a^*_{DA,\lambda}$) and the expected-utility-averaging decision-rule ($a^*_{UA,\lambda}$), as well as optimal abatement (left figure) and optimal saving (right figure) with expert 1's beliefs and expert 2's beliefs (in dashed lines).

To gain intuition, let us focus on the climate mitigation example, the precautionary savings case being fully similar. Expert 2 predicts a lower marginal impact of abatement than expert 1 in terms of expected utility. This explains why a^*_2 is lower than a^*_1. Besides, expert 2 predicts a lower expected utility than expert 1. Indeed, expert 2 predicts more risky damages, which generates a lower expected utility. Thus an increase in abatement widens the distance between the expected utility levels computed with the two expert beliefs. This finally implies that, with the expected-utility-averaging decision-rules, an increase of ambiguity aversion leads to a decrease of the optimal abatement level.

Remark that as a consequence of the Pareto condition, the expected-utility-averaging decision-rules provide an optimal action $a^*_{UA,\lambda}$ that always remains in between a^*_2 and a^*_1. By contrast, distribution-averaging decision-rules can entail abatement or saving levels $a^*_{DA,\lambda}$ that are larger than both a^*_1 and a^*_2. This actually occurs when disagreement aversion is strong enough. Naturally, one may wonder whether it can be rational to abate or save more than what both experts suggest. If the decision-maker could be sure that one expert is right, without knowing necessarily which one, such a decision would hardly be tenable, as choosing action a_1^* would yield a higher utility for sure. However, is there any reason to think that at least one expert is right? A decision-maker who observes that experts disagree may conclude that there is some uncertainty in scientific knowledge and that most likely neither of the experts is completely right. The underlying principle of the distribution-averaging decision-rule is that the decision-maker uses the diversity of expert opinions to infer "the range of possibilities", and then decides to form her own belief using an aggregation procedure which is deliberately cautious in order to reflect disagreement aversion.

3 Conclusion

Decision-makers rely on scientific expertise. It is only natural, and even a sign of healthy science, that experts disagree. But how should decision-makers deal with expert disagreement? Most prominent is the ambiguity neutral approach: The decision-maker forms a (weighted) average of expert views. While consistent and well-understood, the neutral approach has been criticized for implying that a situation of scientific consensus and a situation of wide expert disagreement around the same average belief are treated the same.

In this paper, we have developed a model of cautious aggregation of beliefs, exhibiting a novel notion of disagreement aversion. In line with existing contributions that deviate from the neutral approach, we borrow from the literature on ambiguity aversion. But instead of suggesting a pessimistic aggregation of utility values, as is standard in the literature, our model rests on the pessimistic aggregation of expert beliefs. We demonstrate that our model is more sensitive to disagreement than common models: It allows aversion to "spurious unanimity", situations in which experts agree on the action to be taken but offer different reasons for their assessment.

Our framework, which can be viewed as reflecting a form of the precautionary principle, can readily be applied to concrete policy issues. We have shown that the framework produces intuitive results in the sense that stronger disagreement aversion induces more cautious decisions. This is an interesting feature, which contrasts with what is found with Paretian models, according to which an increase in ambiguity aversion may lead to take less cautious choices. Our work moreover has relevance for the long-standing discussion about the appropriate institutional interplay of scientific expertise and policy-makers, in particular the interplay of risk assessment and risk management (National Research Council 2009). Our contribution strengthens the view that the realms of expertise and decision should be kept separate. According to our approach, the fact that experts reach a consensus about which action to take does not imply that this action should necessarily be chosen. In practice, this means that decisions should not be fully delegated to expert committees. Instead, each expert should submit her scientific assessment. Observing why experts' assessments differ is valuable for a cautious decision-maker, even if the experts agree on the decision.

References

Aczél, János. 1966. *Lectures on functional equations and their applications.* Academic Press. 41

Alon, Shiri, and Gabi Gayer. 2016. "Utilitarian preferences with multiple priors." *Econometrica*, 84(3): 1181–1201. 7

Anscombe, Francis J., and Robert J. Aumann. 1963. "A definition of subjective probability." *Annals of Mathematical Statistics*, 34(1): 199–205. 8

Baillon, Aurélien, Laure Cabantous, and Peter P. Wakker. 2012. "Aggregating imprecise or conflicting beliefs: An experimental investigation using modern ambiguity theories." *Journal of Risk and Uncertainty*, 44(2): 115–147. 3

Barrieu, Pauline, and Bernard Sinclair-Desgagné. 2006. "On precautionary policies." *Management Science*, 52(8): 1145–1154. 7

Basili, Marcello, and Alain Chateauneuf. 2020. "Aggregation of experts' opinions and conditional consensus opinion by the Steiner point." *International Journal of Approximate Reasoning*, 123: 17–25. 6

Berger, Loïc. 2014. "Precautionary saving and the notion of ambiguity prudence." *Economics Letters*, 123(2): 248–251. 28

Berger, Loïc, Johannes Emmerling, and Massimo Tavoni. 2016. "Managing catastrophic climate risks under model uncertainty aversion." *Management Science*, 63(3): 749–765. 6, 28

Berger, Loïc, Nicolas Berger, Valentina Bosetti, Itzhak Gilboa, Lars Peter Hansen, Christopher Jarvis, Massimo Marinacci, and Richard D. Smith. 2021. "Rational policymaking during a pandemic." *Proceedings of the National Academy of Sciences*, 118(4). 3, 6

Bommier, Antoine. 2017. "A dual approach to ambiguity aversion." *Journal of Mathematical Economics*, 71: 104–118. 5, 6, 15, 18

Brandl, Florian. 2020. "Belief-averaged relative utilitarianism." *arXiv preprint 2005.03693*. 7

Castagnoli, Erio, and Marco LiCalzi. 1996. "Expected utility without utility." *Theory and Decision*, 41(3): 281–301. 11

Cerreia-Vioglio, Simone, Lars Peter Hansen, Fabio Maccheroni, and Massimo Marinacci. 2020. "Making decisions under model misspecification." *University of Chicago, Becker Friedman Institute for Economics Working Paper 2020-103*. 6

Cerreia-Vioglio, Simone, Paolo Ghirardato, Fabio Maccheroni, Massimo Marinacci, and Marciano Siniscalchi. 2011. "Rational preferences under ambiguity." *Economic Theory*, 48(2): 341–375. 19, 20

Chateauneuf, Alain. 1999. "Comonotonicity axioms and rank-dependent expected utility theory for arbitrary consequences." *Journal of Mathematical Economics*, 32(1): 21–45. 17

Crès, Hervé, Itzhak Gilboa, and Nicolas Vieille. 2011. "Aggregation of multiple prior opinions." *Journal of Economic Theory*, 146(6): 2563–2582. 3, 6

Danan, Eric, Thibault Gajdos, Brian Hill, and Jean-Marc Tallon. 2016. "Robust social decisions." *American Economic Review*, 106(9): 2407–2425. 7

DeGroot, Morris H., and Julia Mortera. 1991. "Optimal linear opinion pools." *Management Science*, 37(5): 546–558. 6

Gajdos, Thibault, and Jean-Christophe Vergnaud. 2013. "Decisions with conflicting and imprecise information." *Social Choice and Welfare*, 41(2): 427–452. 3, 6

Gajdos, Thibault, Takashi Hayashi, Jean-Marc Tallon, and Jean-Christophe Vergnaud. 2008. "Attitude toward imprecise information." *Journal of Economic Theory*, 140(1): 27–65. 3, 4

Genest, Christian, and James V. Zidek. 1986. "Combining probability distributions: A critique and an annotated bibliography." *Statistical Science*, 1(1): 114–135. 6

Ghirardato, Paolo, and Massimo Marinacci. 2002. "Ambiguity made precise: A comparative foundation." *Journal of Economic Theory*, 102(2): 251–289. 4, 21

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci. 2004. "Differentiating ambiguity and ambiguity attitude." *Journal of Economic Theory*, 118(2): 133–173. 22

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin expected utility with non-unique prior." *Journal of Mathematical Economics*, 18(2): 141–153. 11

Gilboa, Itzhak, Dov Samet, and David Schmeidler. 2004. "Utilitarian aggregation of beliefs and tastes." *Journal of Political Economy*, 112(4): 932–938. 3, 7

Gilboa, Itzhak, Larry Samuelson, and David Schmeidler. 2014. "No-betting-Pareto dominance." *Econometrica*, 82(4): 1405–1442. 7

Gollier, Christian. 2011. "Portfolio choices and asset prices: The comparative statics of ambiguity aversion." *Review of Economic Studies*, 78(4): 1329–1344. 28

Gollier, Christian, and Nicolas Treich. 2003. "Decision-making under scientific uncertainty: The economics of the precautionary principle." *Journal of Risk and Uncertainty*, 27(1): 77–103. 7

Gollier, Christian, Bruno Jullien, and Nicolas Treich. 2000. "Scientific progress and irreversibility: An economic interpretation of the 'Precautionary Principle'." *Journal of Public Economics*, 75(2): 229–253. 7

Hansen, Lars Peter, and Thomas J. Sargent. 2001. "Robust control and model uncertainty." *American Economic Review*, 91(2): 60–66. 6

Harsanyi, John C. 1955. "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility." *Journal of Political Economy*, 63(4): 309–21. 7

Heal, Geoffrey, and Antony Millner. 2018. "Uncertainty and ambiguity in environmental economics: Conceptual issues." *Handbook of Environmental Economics*, 4: 439–468. 6

Hill, Brian. 2012. "Unanimity and the aggregation of multiple prior opinions." *Working paper*. 3

Hylland, Aanund, and Richard Zeckhauser. 1979. "The impossibility of Bayesian group decision making with separate aggregation of beliefs and values." *Econometrica*, 47(6): 1321–1336. 7

Jose, Victor Richmond R., Yael Grushka-Cockayne, and Kenneth C. Lichtendahl. 2013. "Trimmed opinion pools and the crowd's calibration problem." *Management Science*, 60(2): 463–475. 6

Kihlstrom, Richard E, and Leonard J Mirman. 1974. "Risk aversion with many commodities." *Journal of Economic Theory*, 8(3): 361–388. 14

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. "A smooth model of decision making under ambiguity." *Econometrica*, 73(6): 1849–1892. 10

Larrick, Richard P., and Jack B. Soll. 2006. "Intuitions about combining opinions: Misappreciation of the averaging principle." *Management Science*, 52(1): 111–127. 6

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. "Ambiguity aversion, robustness, and the variational representation of preferences." *Econometrica*, 74(6): 1447–1498. 11

Machina, Mark J., and Marciano Siniscalchi. 2014. "Ambiguity and ambiguity aversion." In *Handbook of the Economics of Risk and Uncertainty*. Vol. 13, , ed. Mark Machina and Kip Viscusi, 729–807. North-Holland. 3, 6, 15, 20

McConway, Kevin J. 1981. "Marginalization and linear opinion pools." *Journal of the American Statistical Association*, 76(374): 410–414. 6

Meinshausen, Malte, Nicolai Meinshausen, William Hare, Sarah C. B. Raper, Katja Frieler, Reto Knutti, David J. Frame, and Myles R. Allen. 2009. "Greenhouse-gas emission targets for limiting global warming to 2°C." *Nature*, 458(7242): 1158–1162. 25

Millner, Antony, Simon Dietz, and Geoffrey Heal. 2013. "Scientific ambiguity and climate policy." *Environmental and Resource Economics*, 55(1): 21–46. 6, 28

Mongin, Philippe. 1995. "Consistent Bayesian aggregation." *Journal of Economic Theory*, 66(2): 313–351. 3, 7

Mongin, Philippe. 1998. "The paradox of the Bayesian experts and state-dependent utility theory." *Journal of Mathematical Economics*, 29(3): 331–361. 3

Mongin, Philippe. 2016. "Spurious unanimity and the Pareto principle." *Economics and Philosophy*, 32(3): 511–532. 3, 4, 7

Morris, Peter A. 1977. "Combining expert judgments: A Bayesian approach." *Management Science*, 23(7): 679–693. 6

Nascimento, Leandro. 2012. "The ex ante aggregation of opinions under uncertainty: Aggregation of opinions under uncertainty." *Theoretical Economics*, 7(3): 535–570. 3

National Research Council. 2009. "Science and decisions: Advancing risk assessment." National Academies Press, Washington, D.C. 33

Nau, Robert F. 2006. "Uncertainty aversion with second-order utilities and probabilities." *Management Science*, 52(1): 136–145. 10

Peter, Richard. 2019. "Revisiting precautionary saving under ambiguity." *Economics Letters*, 174: 123–127. 28

Quiggin, John. 1982. "A theory of anticipated utility." *Journal of Economic Behavior and Organization*, 3(4): 323–343. 11

Savage, Leonard J. 1954. *The foundations of statistics.* Courier Corporation. 8

Siniscalchi, Marciano. 2009. "Vector expected utility and attitudes toward variation." *Econometrica*, 77(3): 801–855. 11

Skiadas, Costis. 2013. "Scale-invariant uncertainty-averse preferences and source-dependent constant relative risk aversion." *Theoretical Economics*, 8(1): 59–93. 3

Stanca, Lorenzo. 2021. "Smooth aggregation of Bayesian experts." *Journal of Economic Theory*, 196. 3, 7

Stone, Mervyn. 1961. "The opinion pool." *Annals of Mathematical Statistics*, 32(4): 1339–1342. 6

Strzalecki, Tomasz. 2013. "Temporal resolution of uncertainty and recursive models of ambiguity aversion." *Econometrica*, 81(3): 1039–1074. 15

Yaari, Menahem E. 1969. "Some remarks on measures of risk aversion and on their uses." *Journal of Economic Theory*, 1(3): 315–329. 13

A Appendix

A.1 Proof of Proposition 1

For any expertise $\mathscr{P}$, notice that to any prospect α, we can associate a unique vector $(\sigma_1,\ldots,\sigma_{K_\alpha})$ (defined as in Definition 5) and a unique matrix $\left(q_k^i\right)_{i,k}$ of transformed probabilities ($q_k^i = f\left(p_k^i\right), \forall i,k$) as in the Proposition's statement. Reciprocally, to any such pair $(\sigma_k)_k$, $\left(q_k^i\right)_{i,k}$, we can construct a corresponding prospect α, even if it means choosing a suitable expertise $\mathscr{P} = (P_1,\ldots,P_N)$. Consider either a prospect α or its corresponding pair $(\sigma_k)_k$, $\left(q_k^i\right)_{i,k}$. As the decision-rule is RDU on distribution-consensual prospects, if α is distribution-consensual we have $U^{P_i}(\alpha) = U^{\mathscr{P}}(\alpha), \forall i$. Thus $\sum_{k=1}^{K_\alpha} \sigma_k I\left(q_k^1,\ldots,q_k^N\right) = U^{\mathscr{P}}(\alpha) = \max_i U^{P_i}(\alpha) = \max_{1\leq i\leq N} \sum_{k=1}^{K_\alpha} \sigma_k q_k^i$. Now, notice that α is non distribution-consensual if and only if there are indices $i,j,k.$ such that $q_k^i \neq q_k^j$ and $\sigma_k > 0$. Take any such $(\sigma_k)_k$, $\left(q_k^i\right)_{i,k}$ and consider the outcome x such that $u(x) = \max_i U^{P_i}(\alpha)$, so that $x \succsim^{P_i} \alpha$, $\forall i$. Then, if the decision-rule is disagreement-averse, we have $x \succ^{\mathscr{P}} \alpha$, i.e. $\max_i \sum_{k=1}^{K_\alpha} \sigma_k q_k^i > \sum_{k=1}^{K_\alpha} \sigma_k I\left(q_k^1,\ldots,q_k^N\right)$. Reciprocally, take any non-distribution-consensual prospect α and consider any outcome x such that $x \succsim^{P_i} \alpha$, $\forall i$. If the Proposition's condition holds, we have $U^{\mathscr{P}}(\alpha) < \max_i U^{P_i}(\alpha) \leq u(x)$, i.e. $x \succ^{\mathscr{P}} \alpha$. The decision-rule is therefore disagreement averse.

A.2 Proof of Proposition 2

For the first condition:

$$I\left(q_k^1,\ldots,q_k^N\right) \leq \sum_{i=1}^{N} \lambda_i q_k^i \Rightarrow \sum_{k=1}^{K} \sigma_k I\left(q_k^1,\ldots,q_k^N\right) \leq \sum_{i=1}^{N} \lambda_i \sum_{k=1}^{K} \sigma_k q_k^i \leq \max_{1\leq i\leq N} \sum_{k=1}^{K} \sigma_k q_k^i.$$

For the second condition: Set $\sigma_{K+1} = 1 - \sum_{k=1}^{K} \sigma_k$ and $q_{K+1}^i = 0$ for all $1 \leq i \leq N$. Then, using successively that I is concave and increasing:

$$\sum_{k=1}^{K+1} \sigma_k I\left(q_k^1,\ldots,q_k^N\right) \leq I\left(\sum_{k=1}^{K+1} \sigma_k q_k^1,\ldots,\sum_{k=1}^{K+1} \sigma_k q_k^N\right) \leq \max_{1\leq i\leq N} \sum_{k=1}^{K} \sigma_k q_k^i.$$

In both cases the first inequality is strict when one has $q_k^i \neq q_k^j$ and $\sigma_k > 0$ for some indices i, j, k.

A.3 Proof of Propositions 3 and 5

We first prove Proposition 3 and then explain how the proof can be adapted to prove Proposition 5.

Proof of Proposition 3.

$\Rightarrow$ Assume that $\succcurlyeq_A$ exhibits greater disagreement aversion than $\succcurlyeq_B$. Then, using the reciprocal of the first implication and the second implication in Definition 3, for any expertise $\mathscr{P}$, and any pair (α, β) of distribution-consensual prospects, one has $\alpha \succcurlyeq_A^{\mathscr{P}} \beta \iff \alpha \succcurlyeq_B^{\mathscr{P}} \beta$. Thus, $\succcurlyeq_A$ and $\succcurlyeq_B$ share the same RDU representation on distribution-consensual prospects, and we shall denote u and f their common utility index and probability transformation.

Now, take any expertise $\mathscr{P}$ and any non-constant vector $\vec{p} = (p_1, \ldots, p_N)$. Let $\vec{q} = (f(p_1), \ldots, f(p_N))$. Define $x = u^{-1}(I_A(\vec{q}))$. Denote by $\left(\vec{X}, \vec{p}\right)$ the prospect α with only extremal outcomes (i.e. $\alpha_k \in \{X^-, X^+\}, \forall k$) and such that $D_\alpha^{P_i}(X^+) = p_i, \forall i$. By comparative disagreement aversion, $\left(\vec{X}, \vec{p}\right) \sim_A^{\mathscr{P}} x \Rightarrow \left(\vec{X}, \vec{p}\right) \succ_B^{\mathscr{P}} x$. By definition, $\left(\vec{X}, \vec{p}\right) \sim_A^{\mathscr{P}} x$ if and only if $U_A\left(\left(\vec{X}, \vec{p}\right)\right) = u(x) = I_A(\vec{q})$, which holds by assumption. Thus, $\left(\vec{X}, \vec{p}\right) \succ_B^{\mathscr{P}} x$, i.e. $I_A(\vec{p}) < I_B(\vec{p})$.

$\Leftarrow$ Take any non distribution-consensual prospect α, distribution-consensual prospect β, and expertise $\mathscr{P}$ such that $\alpha \sim_A^{\mathscr{P}} \beta$. Defining $\left(p_k^i\right)$ the probabilities associated to α, notice that $\vec{p}_k = \left(p_k^1, \ldots, p_k^N\right)$ is non-constant for some k so that $I_A(\vec{p}_k) < I_B(\vec{p}_k)$ for those k. For the remaining k where $\vec{p}_k = (p_k, \ldots, p_k)$ is constant, $I_A(\vec{p}_k) = I_B(\vec{p}_k) = p_k$. As $\succcurlyeq_A$ and $\succcurlyeq_B$ share the representation functions f and u and by Definition 5 or 6, this implies $U_B^{\mathscr{P}}(\beta) = U_A^{\mathscr{P}}(\beta) = U_A^{\mathscr{P}}(\alpha) < U_B^{\mathscr{P}}(\alpha)$, i.e. $\alpha \succ_B^{\mathscr{P}} \beta$.

Proof of Proposition 5. The $\Rightarrow$ direction can be proven just as above. For $\Leftarrow$, i.e. to prove the comparative ambiguity aversion, one needs to cover the three cases of Definition 7. To prove the last case, where $\alpha \sim_A \beta$ is non-utility-consensual, one just needs to adjust the $\Leftarrow$ part of the above proof, by taking a non-utility consensual (instead of non-distribution consensual) α. To prove the other cases, it suffices to use that $I_A \leq I_B$ in the representations U_A and U_B.

A.4 Link between the Pareto condition and linear pooling

We give here only the sketch of the proof of Proposition 4, which states that a decision-rule is both expected-utility-averaging and distribution-averaging if and only if it is a linear pooling. Indeed, we prove rigorously a stronger result in online Appendix OA.1: That a distribution-averaging decision-rule is Paretian if and only if it is a linear pooling. That a linear pooling is all at the same time distribution-averaging, expected-utility-averaging, and hence Paretian, can be seen from the Definitions. Now, if a decision-rule is both expected-utility-averaging and distribution-averaging, it allows two representations, which are equal up to an increasing bijection. Considering distribution-consensual prospects, and then looking successively at prospects with only one outcome, with only extremal outcomes, and with two outcomes (one of which we vary), it can be shown that the bijection is indeed the identity function, and that the two representations share the same functions u and f. Considering non-distribution-consensual prospects with only extremal outcomes, it follows that the two representations also share the averager I. Considering prospects such that $\sigma_k = \frac{1}{K_\alpha}, \forall k$, we thus obtain the functional equation for I:

$$\sum_{k=1}^{K_\alpha} \frac{1}{K_\alpha} I\left(f\left(p_k^1\right), \ldots, f\left(p_k^N\right)\right) = I\left(\sum_{k=1}^{K_\alpha} \frac{1}{K_\alpha} f\left(p_k^1\right), \ldots, \sum_{k=1}^{K_\alpha} \frac{1}{K_\alpha} f\left(p_k^N\right)\right),$$

which holds on all $\left(p_k^i\right)$ such that $p_1^i \geq \ldots \geq p_{K_\alpha}^i, \forall i$. Without the latter restriction on the domain, this would exactly be Jensen's functional equation, whose solution is known to be affine (e.g., Aczél 1966). To handle the domain restriction, one can notice that the solution applies locally to any neighborhood in the interior of the domain, and use the connectedness of the domain to show that the affine function is the same on all these neighborhoods. Given that $I(0, \ldots, 0) = 0$, the averager I is here linear.

A.5 Proof of Proposition 6

Take any expertise $\mathscr{P}$, prospect α with at least one non-extremal outcome, and distribution-consensual prospect β. As $\succsim_{DA}$ and $\succsim_{UA}$ share the functions u and f, they coincide on distribution-consensual prospects, and we can denote $U^{\mathscr{P}}(\beta) := U^{\mathscr{P}}_{UA}(\beta) = U^{\mathscr{P}}_{DA}(\beta)$. Also, if α is distribution-consensual, $\alpha \sim^{\mathscr{P}}_{DA} \beta \Rightarrow \alpha \sim^{\mathscr{P}}_{UA} \beta$. For the remainder of the proof, we consider the case where α is non distribution-consensual, so that $\left(p^1_k, \ldots, p^N_k\right)$ is non-constant for some k such that $\sigma_k \neq 0$. Set $\sigma_{K_\alpha+1} = 1 - \sum_{k=1}^{K_\alpha} \sigma_k$ and $p^i_{K_\alpha+1} = 0$ for all $1 \leq i \leq N$. The strict concavity inequality yields: $\sum_{k=1}^{K_\alpha+1} \sigma_k I\left(f\left(p^1_k\right), \ldots, f\left(p^N_k\right)\right) < I\left(\sum_{k=1}^{K_\alpha+1} \sigma_k f\left(p^1_k\right), \ldots, \sum_{k=1}^{K_\alpha+1} \sigma_k f\left(p^N_k\right)\right)$. That is $U^{\mathscr{P}}_{DA}(\alpha) < U^{\mathscr{P}}_{UA}(\alpha)$. Thus, $\alpha \succsim^{\mathscr{P}}_{DA} \beta \Rightarrow U^{\mathscr{P}}_{UA}(\alpha) > U^{\mathscr{P}}_{DA}(\alpha) \geq U^{\mathscr{P}}(\beta) \Rightarrow \alpha \succ^{\mathscr{P}}_{UA} \beta$. To conclude, notice that we have covered all cases needed to prove that $\succsim_{DA}$ exhibits greater ambiguity aversion and greater disagreement aversion than $\succsim_{UA}$.

A.6 Proof of Proposition 7

Set $U_\tau(a) = \sum_{k=1}^{K} \sigma_k I_\tau\left(f\left(p^1_k\right), \ldots, f\left(p^N_k\right)\right)$. The first-order condition of (9) is:

$$U'_\tau(a) = \sum_{k=2}^{K} \left(\partial_1 u(a, \omega_k) - \partial_1 u(a, \omega_{k-1})\right) I_\tau\left(f\left(p^1_k\right), \ldots, f\left(p^N_k\right)\right) + \partial_1 u(a, \omega_1) = 0, \tag{12}$$

given that we have $I_\tau\left(f\left(p^1_1\right), \ldots, f\left(p^N_1\right)\right) = I_\tau(1, \ldots, 1) = 1$. Since decision-maker A is more disagreement-averse than B, we have $I_A\left(f\left(p^1_k\right), \ldots, f\left(p^N_k\right)\right) \leq I_B\left(f\left(p^1_k\right), \ldots, f\left(p^N_k\right)\right)$ for all k with strict inequality for some k since the group of experts disagrees. If $\partial_1 u(a, \omega)$ strictly increases with ω, we have $U'_A(a) < U'_B(a)$. The function $u(a, \omega)$ being strictly concave in a one has $U''_A < 0$ and $U''_B < 0$. Thus, it must be the case that $a^*_A < a^*_B$. If $\partial_1 u(a, \omega)$ strictly decreases with ω, we have $U'_A(a) > U'_B(a)$ and $a^*_A > a^*_B$.

A.7 Proof of Proposition 8

Consider the smooth "KMM" form (1) with functions ψ_τ. Set $V_\tau(a) = \sum_{i=1}^{N} \lambda_i \psi_\tau\left(\sum_{k=1}^{K} \sigma_k f\left(p^i_k\right)\right)$. Solving (10) is equivalent to maximizing $V_\tau(a)$

which gives the first-order condition:

$$V_\tau'(a) = \sum_{i=1}^{N} \lambda_i \psi_\tau' \left(\sum_{k=1}^{K} \sigma_k f\left(p_k^i\right) \right) \cdot \left(\sum_{k=2}^{K} \left(\partial_1 u(a,\omega_k) - \partial_1 u(a,\omega_{k-1})\right) f\left(p_k^i\right) + \partial_1 u(a,\omega_1) \right) = 0, \tag{13}$$

which rewrites:

$$\sum_{i=1}^{N} \underbrace{\frac{\lambda_i \psi_\tau' \left(\sum_{k=1}^{K} \sigma_k f\left(p_k^i\right)\right)}{\sum_{l=1}^{N} \lambda_l \psi_\tau' \left(\sum_{k=1}^{K} \sigma_k f\left(p_k^l\right)\right)}}_{\tilde{\lambda}_i(\psi_\tau,a)} \cdot \underbrace{\left(\sum_{k=2}^{K} \left(\partial_1 u(a,\omega_k) - \partial_1 u(a,\omega_{k-1})\right) f\left(p_k^i\right) + \partial_1 u(a,\omega_1) \right)}_{\rho_i(a)} = 0. \tag{14}$$

Since the experts are ordered in the sense of first-order stochastic dominance and $\partial_2 u(a,\omega) > 0$, we have $\sum_{k=1}^{K} \sigma_k f\left(p_k^1\right) \leq \sum_{k=1}^{K} \sigma_k f\left(p_k^2\right) \leq \ldots \leq \sum_{k=1}^{K} \sigma_k f\left(p_k^N\right)$, with at least one strict inequality given that the group of experts disagrees. In (14), we can view $\tilde{\lambda}_i(\psi_\tau,a)$ as a distribution function where i would be the random variable. By hypothesis, we have $\psi_A = h \circ \psi_B$ where h is an increasing and (strictly) concave function. The likelihood ratio of $\tilde{\lambda}_i(\psi_A,a)$ and $\tilde{\lambda}_i(\psi_B,a)$ then writes:

$$\frac{\tilde{\lambda}_i(\psi_A,a)}{\tilde{\lambda}_i(\psi_B,a)} = h' \left(\psi_B \left(\sum_{k=1}^{K} \sigma_k f\left(p_k^i\right) \right) \right) \cdot \frac{\sum_{l=1}^{N} \lambda_l \psi_B' \left(\sum_{k=1}^{K} \sigma_k f\left(p_k^l\right)\right)}{\sum_{l=1}^{N} \lambda_l \psi_A' \left(\sum_{k=1}^{K} \sigma_k f\left(p_k^l\right)\right)}. \tag{15}$$

Since we have $\sum_{k=1}^{K} \sigma_k f\left(p_k^i\right)$ increasing with i with at least one strict inequality, $\psi_B' > 0$, $\psi_A' > 0$ and $h'' < 0$, the likelihood ratio (15) is decreasing with i with at least one strict inequality. Thus, the distribution $\tilde{\lambda}_i(\psi_B,a)$ strictly dominates the distribution $\tilde{\lambda}_i(\psi_A,a)$ in the sense of monotone likelihood ratio, which implies that the former strictly first-order stochastically dominates the latter. If $\partial_1 u(a,\omega)$ strictly increases (decreases) with ω, we have $\rho_1(a) \leq \rho_2(a) \leq \ldots \leq \rho_N(a)$ ($\rho_1(a) \geq \rho_2(a) \geq \ldots \geq \rho_N(a)$) with at least one strict inequality, since the experts are ordered in the sense of first-order stochastic dominance and the group of experts disagrees. As a consequence, since $\tilde{\lambda}_i(\psi_B,a)$ strictly first-order stochastically dominates $\tilde{\lambda}_i(\psi_A,a)$, we get, for a given a, $\sum_{i=1}^{N} \tilde{\lambda}_i(\psi_A,a)\rho_i(a) < \sum_{i=1}^{N} \tilde{\lambda}_i(\psi_B,a)\rho_i(a)$ ($\sum_{i=1}^{N} \tilde{\lambda}_i(\psi_A,a)\rho_i(a) > \sum_{i=1}^{N} \tilde{\lambda}_i(\psi_B,a)\rho_i(a)$). Thus, with $a = a_B^*$, we have $\sum_{i=1}^{N} \tilde{\lambda}_i(\psi_A,a_B^*)\rho_i(a_B^*) < 0$ ($\sum_{i=1}^{N} \tilde{\lambda}_i(\psi_A,a_B^*)\rho_i(a_B^*) > 0$), and $V_A'(a_B^*) < 0$

$(V_A'(a_B^*) > 0)$. We know that $V_A'' < 0$, since $u(a,\omega)$ is strictly concave in a and ψ_A concave. Thus, we can conclude that $a_A^* < a_B^*$ $(a_A^* > a_B^*)$.

A.8 Proof of Proposition 9

A direct application of Proposition 7 shows that the optimal action $a^*_{DA,\lambda}$ is increasing with λ. On the other hand, we cannot apply Proposition 8 since experts are not ordered in terms of first-order stochastic dominance. Following a reasoning similar to the proof of Proposition 8, we show that the optimal action $a^*_{UA,\lambda}$ is decreasing with λ. We show this result for the precautionary savings example. The reasoning is the same for the climate mitigation example. With a expected-utility-averaging decision-rule and a smooth "KMM" form, the optimal saving level satisfies (14) with $K = 3$ states of nature, $\omega_1 = 0$, $\omega_2 = 1$, $\omega_3 = 3$, $N = 2$ experts, $p_1^1 = 1$, $p_2^1 = 1$, $p_3^1 = 0$, $p_1^2 = 1$, $p_2^2 = 0.4$, $p_3^2 = 0.4$, $a \in [0, 0.9]$ and $u(a,\omega) = 2 - a - \frac{1}{8}(2-a)^2 + \omega + a - \frac{1}{8}(\omega+a)^2$. In contrast with the proof of Proposition 8, we have $\sum_{k=1}^K \sigma_k f\left(p_k^1\right) = 2 - a - \frac{1}{8}(2-a)^2 + 1 + a - \frac{1}{8}(1+a)^2$ strictly larger than $\sum_{k=1}^K \sigma_k f\left(p_k^2\right) = 2 - a - \frac{1}{8}(2-a)^2 + \frac{12}{10} + a - \frac{6}{80}a^2 - \frac{4}{80}(3+a)^2$ since $\sum_{k=1}^K \sigma_k f\left(p_k^1\right) - \sum_{k=1}^K \sigma_k f\left(p_k^2\right) = \frac{1}{8} + \frac{1}{20}a > 0$ for $a \geq 0$. Considering ψ_A more concave than ψ_B, this implies that the likelihood ratio of $\tilde{\lambda}_i(\psi_A, a)$ and $\tilde{\lambda}_i(\psi_B, a)$ is strictly increasing with i, and the distribution $\tilde{\lambda}_i(\psi_B, a)$ is strictly first-order stochastically dominated by the distribution $\tilde{\lambda}_i(\psi_A, a)$. Moreover, we have $\partial_{1,2}u(a,\omega) = -\frac{1}{4} < 0$ and $\rho_1(a) = -\frac{1}{2}a + \frac{1}{4} > \rho_2(a) = -\frac{1}{2}a + \frac{1}{5}$. Thus, for a given a, $\sum_{i=1}^N \tilde{\lambda}_i(\psi_A, a)\rho_i(a) < \sum_{i=1}^N \tilde{\lambda}_i(\psi_B, a)\rho_i(a)$. This implies $V_A'(a_B^*) < 0$ and $a_A^* < a_B^*$. Finally, this shows that with a KMM kind of averager of the form (11), the optimal saving level $a^*_{UA,\lambda}$ is decreasing with the ambiguity aversion parameter λ.