

The principal-agent problem with smooth ambiguity

Christian Kellner¹ 

Received: 15 October 2015 / Accepted: 31 January 2017 / Published online: 11 February 2017
© The Author(s) 2017. This article is published with open access at Springerlink.com

Abstract We study a principal-agent model in which the (effort-dependent) realisation of output levels is ambiguous, and the agent is ambiguity averse (while the principal is ambiguity neutral). We show that introducing ambiguity aversion will lower profits if the action that the principal wants to implement is the most ambiguous one, while they may increase otherwise. Regarding the design of the optimal contract, we show that under ambiguity aversion the optimal incentive scheme may not be monotone even if a natural generalization of the monotone likelihood ratio property is satisfied, and illustrate how this fact could affect the design of contracts in an applied economic context. We also find that the individual rationality constraint need not bind in the presence of ambiguity aversion unless preferences satisfy constant absolute ambiguity aversion.

Keywords Moral hazard · Ambiguity aversion · Optimal contract

JEL Classification D81 · D86 · M52

1 Introduction

In its most standard form the principal-agent problem is typically described in the following way. The owner of a firm, the principal, has to decide how to remunerate

A substantial part of the work on this paper has been done while the author was affiliated to the Kellogg School of Management at Northwestern University and to the University of Bonn. The author thanks Paolo Ghiradato, Peter Klibanoff, and Marciano Siniscalchi for valuable discussions.

✉ Christian Kellner
c.kellner@soton.ac.uk

¹ University of Southampton, University Road, Southampton SO17 1BJ, UK

a manager, the agent, who is in charge of running the firm. The effort of the agent determines, to a large extent, the quantity produced by the firm, but the firm's output is also influenced by random events that are beyond the control of the manager. The problem arises as the manager's effort cannot be observed (or otherwise inferred) by the principal.

As it is commonly assumed that the principal is risk neutral and the agent risk averse, the solution to this problem arises as the optimal trade-off between risk-sharing and incentives. Thus it should be apparent that the optimal wage scheme is sensitive to the agent's reaction to non-deterministic payoffs. Most of the literature assumes that the agent (and the principal) are expected utility maximisers. In the context of the principal-agent problem, the expected utility framework requires the agent to treat the uncertainty about her monetary payoffs in the same way as a lottery over monetary prizes with known probabilities. The well-known Ellsberg paradox however has convincingly cast some doubt at the validity of such an assumption (Ellsberg 1961). In response, the decision theory literature has generalised the expected utility model to accommodate aversion (as well as other attitudes) towards ambiguity, which can explain the behaviour observed in the Ellsberg paradox. Specifically, ambiguity is defined as subjective uncertainty about outcome distributions. Our paper addresses the question to what extent the traditional analysis of the principal-agent problem remains valid in situations where ambiguity about the consequences of the agent's actions prevails.

Why might ambiguity be of particular interest in the principal-agent problem? In many cases where moral hazard arises in economic interactions, it seems reasonable to assume that the consequences of the agent's actions cannot be described with great confidence by a single probability distribution. A good example may be a firm that employs a scientist to develop a new production technique. Yet, in other instances ambiguity may well be less of an issue. If the principal routinely contracts with some agent to do always essentially the same task, then it should be possible to have a precise understanding about the consequences of each of the agent's actions. This might be true for the relationship between a firm and one of her sales agents. Thus, our approach also enables us to address the question how the optimal incentive scheme will differ between such situations. Additionally, ambiguity can vary between the actions available to the agent. Consider again a scientist: For her, exerting a lot of effort to develop an entirely new approach may have very ambiguous benefits. Exerting low effort (by merely adopting an existing technology) may still not lead to a deterministic outcome, but there may be less ambiguity about the outcome distribution. For the case of a "routine" principal-agent situation only the action which is in fact implemented in equilibrium might be well understood. The other actions available to the agents might very well be of ambiguous consequences (as they are never actually chosen).¹

In the principal-agent problem the optimal incentive contract is, in general, inefficient. That is, if the principal could observe the action chosen by the agent she could offer a different wage scheme that makes her better off without harming the agent. As long as the optimal incentive scheme does not leave any rent to the agent, the

¹ This point has been raised first by Ghirardato (1994). The observation that ambiguity might differ between actions (or outcomes) will actually turn out to be of some importance in determining when ambiguity decreases the principal's profit.

degree of inefficiency is directly linked to the principal's profit. Therefore, we want to find out whether profits necessarily decrease when ambiguity matters. Ambiguity may also affect other properties of the optimal incentive scheme. Wage schedules are often expected to be monotonic. This means that if an outcome realises that is better for the principal, the agent receives a higher wage as well. Even without ambiguity additional assumptions are needed to ensure the optimality of monotonic schemes. Thus we seek to show whether similar assumptions are sufficient in the presence of ambiguity, or if stronger assumptions are needed.

To address these questions we adopt a recent model of decision making under ambiguity, the smooth ambiguity model (Klibanoff et al. 2005). This is a key difference between this paper and the work by Ghirardato (1994) (which predates the smooth model), who studies the principal-agent problem with a different model of ambiguity. In his paper, the agents use capacities instead of probabilities, which are evaluated using Schmeidler's Choquet expected utility theory (Schmeidler 1989). A major advantage of using the smooth ambiguity model is that it suggests a clear distinction between ambiguity and attitude towards ambiguity, and allows each to vary separately, while in the Choquet expected utility model these two concepts cannot be separated.² For most parts of the paper, we assume that absolute ambiguity aversion is constant. We will show that without this assumption the individual rationality constraint might not bind in the optimal contract, so that the agent can be strictly better off than under her outside option.

Lang (forthcoming) contains a result complementary to this paper for the case of infinitely many effort levels. He finds that under the smooth ambiguity model (and other second-order models) a strictly positive effort level is typically optimal while under first order models of ambiguity aversion, like the maximin expected utility model (Gilboa and Schmeidler 1989) or the Choquet expected utility model, zero effort may be optimal in some circumstances. Lopomo et al. (2011) show that given Bewley preferences (which can be interpreted as ambiguity averse preferences), the optimal contract may be coarser due to ambiguity. Contracting under vague information is studied in Viero (2012), the effects of loss aversion are investigated in Herweg et al. (2010). Carroll (2015) find that ambiguity about available actions can lead to linear contracts. Di Tillio et al. (2014) and Bose and Renou (2014) provide rationales for designing ambiguous contracts even if there is initially no ambiguity about the relevant distributions. We restrict ourselves instead to studying standard (i.e. unambiguous) contracts in an exogeneously ambiguous setting. For the case of more than one agent, Kellner (2015) shows ambiguity aversion can make the use of tournaments more attractive.

Other papers related to moral hazard and ambiguity are Mukerji (1998, 2003). For other more distantly related papers on the implications of ambiguity in economics see the survey by Mukerji and Tallon (2004). Also of interest, as another application of the smooth ambiguity model, in this case to portfolio choice, is Gollier (2011). While it is beyond the scope of this introduction to give an overview of the contributions to the principal-agent literature, it is worth mentioning that our formulation of the problem

² To achieve this separation between ambiguity and ambiguity aversion the smooth ambiguity model assumes that the agent's preferences over second-order acts (bets on the true probability) can be observed. See Epstein (2010) for a critical discussion.

most directly corresponds to [Grossman and Hart \(1983\)](#) (henceforth GH), which we also use as a reference when we compare our results to the case without ambiguity aversion.

Even in the absence of ambiguity, very few general results regarding properties of the optimal contract are available.³ Hence, we will follow established practice of restricting the model to special cases, most notably the case of two actions and two outcomes, as this allows us to highlight important channels how ambiguity and ambiguity aversion matters.

Regarding the shape of the optimal contract, we will build on the observation of [Ghirardato \(1994\)](#), who finds that, under ambiguity, non-monotonicity may arise even in the case of only two outcomes. We will however also show that in many cases the optimal wage contract will nevertheless entail a positive bonus to the agent if the better outcome realises, for instance if it is in a certain sense unambiguous that the high-cost action results more likely in a better outcome. On the other hand, we identify a new source of non-monotonicity for the case of at least three outcomes: Non-monotonicities may arise if there is more ambiguity about the probability of some outcomes than others. Then the principal might aim at reducing the ambiguity about wages by keeping them similar between the more ambiguous outcomes. We will argue using an example that this effect is not simply a curiosity, but may be quite relevant in the context of real-world incentive contracts.

In addition to describing the optimal wage contract, we provide some insight on how the principal's profits change in response to changes in the problem. In particular, as the smooth ambiguity model provides a separation between ambiguity and ambiguity aversion, we can study an increase in ambiguity and an increase in ambiguity aversion separately. We will show that these two changes may have very different effects. Replacing an ambiguity neutral agent with an ambiguity averse agent, for instance, will be bad for the principal, if she wants to implement the most ambiguous action. If the principal implements the least ambiguous action, then ambiguity aversion might actually increase profits, in marked contrast to the introduction of risk aversion in the standard model. The intuition behind this finding is that, while ambiguity aversion always makes the incentive constraint harder to satisfy, the effect on the individual rationality constraints depends on which action is the more ambiguous one. The situation is different, however, for an increase in ambiguity. Using a suitable definition of a uniform increase in ambiguity, we find that if ambiguity increases uniformly, then the principal's payoff may increase only if actually the most ambiguous action is to be implemented. This follows since, given our specification of preferences, a payment scheme that is based on the most ambiguous action is less affected from a further increase in ambiguity than any less ambiguous payment scheme.

Finally, we look at the case of more than two actions, and show that, given smooth ambiguity aversion, it may be the case that the binding incentive constraint in an optimal contract pertains to a more expensive action. [Table 1](#) summaries the main differences between the expected utility approach, assuming ambiguity neutrality (col-

³ For instance, if there are more than two actions it is unclear which incentive constraint will be the relevant one. Without assuming the monotone likelihood ratio property it is possible that the optimal contract is not monotonic.

Table 1 Properties of the optimal contract

	Ambiguity neutrality	Smooth ambiguity aversion	Choquet EU
IR binding	Yes	No, requires, e.g., CAAA	Yes
Monotonicity (binary outcomes)	Yes	No, unless, e.g., actions ranked unambiguously	No
Monotonicity (many outcomes)	Under MLRP	No	No
Preference change (two outcomes)	Risk aversion typically reduces profits	Ambiguity aversion often increases profits	Ambiguity aversion often increases profits
IC binds	For cheaper action	In either direction	For cheaper action

umn 1), and the model allowing for smooth ambiguity aversion (column 2). Column 3 compares our results with the Choquet model using the results in [Ghirardato \(1994\)](#).

The remainder of this paper is organised as follows: Sect. 2 introduces the model. Section 3 discusses the properties of the optimal contract. It also motivates why we restrict the model to the case of constant absolute ambiguity aversion in the remaining parts of the paper. Section 4 discusses two properties of the optimal contract for the case of two outcomes only, assuming for the larger part that there are only two actions: First, we discuss sufficient conditions for the monotonicity of the optimal contract. Second, we discuss comparative statics in ambiguity and ambiguity aversion. Section 5 discusses first the robustness of these results to the case of more than two outcomes, and second to the case of more than two actions.

2 The model

In the standard formulation of the principal-agent problem, a principal (typically the owner of a firm) has to hire an agent (a worker or manager) to perform a certain task. The agent's effort choice while doing this task cannot be observed directly, but it determines the distribution of a random variable (the output of the firm). Only the principal, not the agent, directly cares about the realisation of this random variable. The principal's problem is to design a payment scheme that incentivises the agent to choose an action that is in the principal's best interest, given the constraint that the scheme cannot depend on the action choice itself.

We will assume that there are only finitely many output levels that the firm might possibly achieve, which are elements of the finite set $Q = \{q_1, \dots, q_I\}$. As a convention we require that $q_i < q_j$ if $i < j$. As the principal can reward the agent based on the output level only, the wage scheme can be represented by a vector $w = (w_1, \dots, w_I)$. The agent can choose from a finite set of actions, A . Choosing an action $a \in A$ incurs an effort cost c^a to the agent. The standard treatment of the principal-agent problem would now assume that in addition to the effort cost, each action a is identified with a probability distribution $p^a \in \Delta(Q)$. That is, when the agent chooses action a , the probability that the firm achieves output level q_i is given by p_i . Our approach

differs here since we introduce ambiguity by assuming that both the principal and the agent have a common but imperfect understanding of the (stochastic) relationship between the actions and the consequences [as axiomatised by [Klibanoff et al. \(2005\)](#)]. Specifically, they are unsure about these probabilities and thus they think that a set of probabilities might better describe the implications of each action. However, they do not necessarily consider all members of that set equally likely, but assign different likelihoods to them, which are specified by a probability measure μ^a defined on $\Delta(Q)$. To rule out any effects that may arise from information asymmetries, we assume that both principal and agent use the same measure to evaluate each act.

To close the description of the agent's preferences we assume, based on the smooth ambiguity model, that they can be represented by the utility function $U : A \times (\underline{w}, \infty)^I \rightarrow \mathbb{R}$ defined by

$$U(a, \mathbf{w}) = \int \phi \left(\sum_{i=1}^I p_i u(w_i) - c^a \right) d\mu^a(p).$$

In this representation, the increasing function $u : (\underline{w}, \infty) \rightarrow \mathbb{R}$ represents the agent's attitude towards risk. We allow \underline{w} to be $-\infty$, but require that $\lim_{w \rightarrow \underline{w}} u(w) = -\infty$. The function $\phi : \mathbb{R} \rightarrow \mathbb{R}$, represents the agent's attitude towards ambiguity. We assume that the agent is ambiguity averse, corresponding to a concave ϕ , and *strictly* risk averse, corresponding to a strictly concave u .⁴ Typically our results carry over to the case where the agent is risk neutral, provided she is *strictly* ambiguity averse.⁵

To interpret the utility representation, observe first that for all actions that can be described by a single, unambiguous probability distribution (i.e. for all actions a where μ^a puts weight 1 on some p^a), the equation specialises to

$$U(a, \mathbf{w}) = \phi \left(\sum_{i=1}^I p_i^a u(w_i) - c^a \right).$$

Thus such actions are ranked by the agent in the same way as by an expected-utility maximizer.

For ambiguous actions, the agent computes the expected utility associated to each possible probability distribution in the same way in a first step. Then she weights the resulting utilities according to her confidence in each probability distribution (given by μ^a), after she has applied the transformation ϕ , which represents her ambiguity attitude. Note that the function $u(a, w) = u(w) - c^a$ determines the agent's preferences between effort costs and wage payments. We have implicitly assumed that this function is additively separable, which is a common simplifying assumption. Note that due to

⁴ We additionally make the standard assumptions that u and ϕ are continuous and strictly increasing.

⁵ There are subtle differences however. It will become clear that the first best could be achieved under risk neutrality despite strict ambiguity aversion whenever there is no payoff-relevant ambiguity about equilibrium wages (for instance, when a unique probability is attributed to some subset of outcomes). Similarly, we will point out that regarding the question whether at least one IC constraint binds, there is an important difference.

the presence of ambiguity this does not necessarily mean however, that the function $U(a, \mathbf{w})$ is additively separable.⁶

Sometimes it is convenient to assume that the support of μ^a is finite for all actions. In this case we write p_{ij} for the probability of realizing output q_i with action a according to the j -th probability distribution and the number μ_j^a denotes the weight assigned to this distribution given action a . Then the agent's utility of choosing action a given wage scheme \mathbf{w} becomes

$$U(a, \mathbf{w}) = \sum_{j=1}^J \mu_j^a \phi \left(\sum_{i=1}^I p_{ij} u(w_i) - c^a \right).$$

Finally, we turn to the principal's payoff. We denote the firm's payoff for action a under payment scheme \mathbf{w} by the function $\Pi : A \times (w, \infty)^I \rightarrow \mathbb{R}$. We assume the firm is both ambiguity and risk neutral (Klibanoff et al. (2005) show that ambiguity neutrality corresponds to ϕ affine). Thus, the firm's payoff can be expressed as

$$\Pi(a, \mathbf{w}) = \int \left(\sum_{i=1}^I p_i (q_i - w_i) \right) d\mu^a = \sum_{i=1}^I \left((q_i - w_i) \int p_i d\mu^a \right).$$

The principal maximises expected net output (i.e. profit), where the expectation is taken with respect to the μ -average distribution $\int p_i d\mu^a$.

The principal's problem As the formulation of the principal's problem differs from Grossman and Hart (1983) only by the way the uncertainty is modelled, we will be brief. It is generally useful to think about the principal's problem in two stages: First, she determines the cost of implementing each of the actions available to the agent. If it is impossible to implement an action, we will say that the cost is infinity. Second, she compares the benefits of each action with the implementation costs to determine which action she wants to implement.

Implementing a given action Suppose that the principal wants to implement action a using reward scheme \mathbf{w} . The agent will work for the principal only if she prefers doing so to the outside option, which corresponds to a certain utility level, $\phi(u_0)$, which does not depend on the output distribution. Moreover, the agent must (weakly) prefer this action over any other action. Thus, the problem of optimally implementing a given action a can be summarized as to maximise $\Pi(a, \mathbf{w})$ by choosing \mathbf{w} subject to the constraints

$$U(a, \mathbf{w}) \geq \phi(u_0) \quad (\text{IR})$$

$$U(a, \mathbf{w}) \geq U(a', \mathbf{w}) \quad \forall a' \in A. \quad (\text{IC})$$

⁶ We could maintain the assumption that $U(a, \mathbf{w})$ is additively separable by choosing $U(a, \mathbf{w}) = \phi^{-1} \left(\int \phi \left(\sum_{i=1}^I p_i u(w_i) \right) d\mu^a(p) \right) - c^a$. Also in this case, the agents would act identical to expected utility decision makers in the absence of ambiguity. Assuming constant absolute ambiguity aversion, the two approaches become equivalent.

It is important to observe that, when the agent is ambiguity neutral (ϕ affine) or there is no ambiguity (μ^a degenerate for each action) this is exactly the standard GH model. The very same arguments as used there establish that this problem always has a solution. (This is formally stated and proven as Lemma 1 in the Appendix.)

3 The optimal incentive scheme

As mentioned in the introduction, even in the standard model it is not generally straightforward to describe the optimal incentive scheme, at least not without making additional assumptions about the model. Yet many interesting properties of the optimal incentive scheme can be stated at least for some special cases of the model. In this section we will present those (limited) results that hold even if we do not impose any constraints on the structure of the problem (that is, we do not limit the number of actions or outcomes), while in the larger part of Sect. 4 we focus on the most important special case (two outcomes and two actions).

The first question we address is whether the agent's IR constraint will bind, or whether she might end up strictly better off than with her outside option. Recall that in the standard model, the IR constraint will bind if utility is either additively (as we have assumed before) or multiplicative separable in monetary reward and effort costs.

Example 1 In this example there are two actions, H and L , with $c^H > c^L$, and two outcomes. Consider a decision maker with ambiguity attitude represented by the function $\phi(u) = -\exp(-\sqrt{u}/2)$. The high-cost action is ambiguous, weight $\mu^H = .7$ is attributed to the distribution where the better outcome obtains for sure ($p_{21}^H = 1$), probability .3 to the probability distribution where the worse outcome is certain, i.e. $p_{22}^H = 0$. Action L unambiguously results in the better outcome with a probability of .5, i.e. $p_{21}^L = p_{22}^L = .5$.⁷ In this example, the individual rationality constraint does not bind in the optimal contract, for reasons illustrated in Fig. 1.

The graph depicts the agent's indifference curves in a state-preference graph, where the axes represent $u(w_1)$ and $u(w_2)$, the interim utilities the agent gets (resulting from some wage scheme) in the two different states. As u is increasing, each point in the graph corresponds to a particular wage scheme (w_1, w_2) .⁸ The curve U_H represents all wage schemes that leave the agent indifferent between not working for the principal (which results in the outside option without any effort costs), and choosing the high-cost action H (after accepting the principal's incentive scheme). In other words, U_H shows all the points $(u(w_1), u(w_2))$ such that the IR constraint binds, $U_H = \{(u(w_1), u(w_2)) : U(H, (w_1, w_2)) = \phi(u_0)\}$, and points to the right satisfy the constraint with a strict inequality. The curve U_L represents all incentive contracts that leave the agent indifferent between choosing the low-cost action (if she accepts the contract) and refusing to work for the principal (which

⁷ To complete the example, assume effort costs are $c^H = .01$ and $c^L = .092$, and $q_2 - q_1$ is very large. The outside option is given by $\phi(u_0) = -1$. A similar example can also be constructed for the case that $U(a, w)$ is additively separable in c^a as suggested in Footnote 6.

⁸ Note that if the agent were risk neutral (i.e. $u(w) = w$) the labels of the axis could be reinterpreted as the dollar amounts the agent receives in the two states.

As the principal is ambiguity neutral, the benefit of implementing action H is directly proportional to $q_2 - q_1$, and hence can be chosen arbitrarily large by choosing a high enough q_2 . Therefore, there is indeed a concrete principal-agent problem where it

 Springer

is optimal to implement the high-cost action, even if this is impossible with a binding IR constraint.¹⁰

Note that such an example cannot be constructed if ambiguity is modeled using the non-additive probabilities framework. Formally, this is established in [Ghirardato \(1994\)](#), Lemma 1, which states that the IR constraint always binds in their model. In a two outcome world it is easy to understand this result as indifference curves are piecewise linear (while they are linear under expected utility), parallel lines with only one kink at the forty-five degree line. So if the indifference curves associated to the two actions don't intersect at the reservation utility, they never intersect. If they do, the intersections corresponding to higher utility levels will be shifted out parallel to the 45 degree line and so they will be always more expensive than the equivalent of the outside option. The same applies to the expected utility framework, where indifference curve are not piecewise linear, but linear.

We will show now that this result still extends to some classes of smooth ambiguity preferences. To do so, we specialise our model to the case of constant absolute ambiguity aversion preferences (CAAA), as defined by [Klibanoff et al. \(2005\)](#). Such preferences correspond to the case where $\phi(u) = -\exp(-\alpha u)$, with α being a parameter measuring the degree of absolute ambiguity aversion.¹¹ The analogy to constant absolute risk aversion is obvious.

Proposition 1 *1. The individual rationality constraint may not bind in an optimal contract.*

2. The individual rationality constraint binds if ϕ is of the CAAA variety.

A proof is given in the Appendix. To understand the intuition, note that only for CAAA preferences the standard proof of GH can be followed: Assume a contract satisfies both constraints, but the IR constraint does not bind. Create a second contract, where utility $u(w_i)$ given any outcome i is lowered by a constant amount. The second contract is actually preferred by the principal. If the amount is small enough, the IR is still satisfied. Moreover, the expected utility under any action (and any possible distribution) decreases by this constant. Hence, under any action, uncertainty about the expected utility remains the same, except that the mean changes by an equal amount. This has no effect on the ranking of the different actions (and consequently the IC constraint) precisely if preferences are of the CAAA variety (including the ambiguity neutral case). Constant ambiguity attitude ensures that the ambiguity premium is translation invariant, which is in general not the case given the smooth model.¹²

In principle a similar argument would hold if *relative* ambiguity aversion (CRAA, $\phi(u) = u^\alpha$) was constant, and the utility function u would be multiplicatively separa-

¹⁰ Note that this example leaves open the question whether the optimal contract would provide the agent with the lowest utility level under which the IR constraint can hold (given that the IC-constraint binds). A different example can be constructed that shows that even that is not true.

¹¹ Note that in this case the agent's preferences could also be represented using other ambiguity models, including variational preferences ([Maccheroni et al. 2006](#)) or the Vector Expected Utility model ([Siniscalchi 2009](#)).

¹² If one drew a graph analogous to Fig. 1 for CAAA preferences, U'^H is a translation of U^H along the 45 degree line using the same vector that makes U'^L a translation of U^L , even if neither curve typically will linear.

ble in costs and wage, i.e. $u(w, a) = c^a u(w)$. However we have seen that to ensure that the IR constraint binds, it is at least necessary that the image of u can never reach a lower bound. But then preserving existence actually requires that also $\lim_{w \rightarrow \underline{w}} u(w) = -\infty$. But this is incompatible with constant relative ambiguity aversion.

In order to facilitate comparison with the standard model we will henceforth assume that preferences are of the CAAA variety. To simplify the notation we will sometimes continue to use $\phi(u)$ even if we have this particular functional form in mind. Moreover we define the payment necessary to obtain a utility level u by the (increasing) function $h(u)$ so that $h(u(w)) = w$. Strict risk aversion implies that h is strictly convex.

We can now state a result for the incentive constraints.

Proposition 2 *Suppose the agent's preferences are CAAA and the principal chooses not to implement the least-cost action. Then at least one of the incentive constraints will bind in any optimal solution.*

The intuition behind this result is the following: Suppose w is a payment scheme where no IC constraint binds. Consider a payment scheme w' where utility levels $u(w'_i)$ are a convex combination to those of w and the reservation utility. The scheme w' still satisfies the IR, as it is less ambiguous, and, if sufficient weight is given to w , all incentive constraints, while it is cheaper for the principal, as it is less risky.

Note that often, but not always, this result would also hold if the agent is risk neutral, but *strictly* ambiguity averse. To see this, assume the optimal contract is not unambiguous (i.e. does not yield the same expected utility under all possible distributions). Since the payments are more similar between outcomes given w' , this scheme also strictly reduces ambiguity about the payments for the agent. Thus, under w' also the IR constraint does not bind, so that the principal can find yet another contract that satisfies all the constraints, but entails a lower ambiguity premium. Hence, unless the optimal contract achieves unambiguous wages, at least one incentive constraint must bind in the optimal contract even under risk neutrality, provided ambiguity aversion is strict. It might, however, well be possible that the distribution of wages in an optimal contract is unambiguous. If, for instance, all possible distributions agree about the probability that one of the highest two outcomes obtains, a contract which pays the same wage level for these two outcomes, and a different wage level for all other outcomes, would result in an unambiguous wage distribution.

We conclude this section by introducing a notation that decomposes the individual's beliefs about each action into two parts. Given our restriction to CAAA preferences it will turn out useful in the following sections to distinguish between the information about each action that is relevant also for an ambiguity neutral agent, and the part that is relevant only to ambiguity averse agents. Thus, we will denote the average probability distribution over output levels by

$$\bar{p}^a \equiv \int p d\mu^a,$$

and the difference between a distribution p^a and the average \bar{p}^a by \hat{p}^a . Hence the ambiguity associated with each action can be described by the measure $\hat{\mu}^a$, where $\hat{\mu}^a(\hat{p}^a) \equiv \mu^a(\hat{p}^a + \bar{p}^a)$ for any $\hat{p}^a \in \text{supp}(\mu^a) - \{\bar{p}^a\}$.

This allows us to restate the problem of implementing action a as

$$\begin{aligned} \min_{\mathbf{u}} \sum_i \bar{p}_i^a h(u_i) \text{ s.t. } \mathbf{u} \in [\text{img}(u)]^I \text{ and} \\ \sum_i \bar{p}_i^a u_i - c^a - \mathcal{F}^a(\mathbf{u}) \geq u_0 \quad (\text{IR}) \\ \sum_i \bar{p}_i^a u_i - c^a - \mathcal{F}^a(\mathbf{u}) \geq \sum_i \bar{p}_i^{a'} u_i - c^{a'} - \mathcal{F}^{a'}(\mathbf{u}) \quad \forall a' \in A, \quad (\text{IC}) \end{aligned}$$

where $\mathcal{F}^a(\mathbf{u}) = \frac{1}{\alpha} \log \int \exp[-\alpha \sum_i \hat{p}_i u_i] d\hat{\mu}^a$. The advantage of this approach is that it highlights the effect of ambiguity aversion. Here $\mathcal{F}^a(\mathbf{u})$ represents the ambiguity premium that the principal needs to pay to the agent. In the individual rationality constraint, the ambiguity premium is the only difference to the standard framework with an unambiguous distribution of \bar{p}^a . For the incentive constraints, the effect of ambiguity aversion corresponds to the difference between the ambiguity premia of the two actions. Via the IR constraint ambiguity aversion will always have a non-positive effect for the principal's profits, the effect via the IC constraint could go either way.

4 Optimal incentives: the two outcome case

To get a first understanding of how the addition of ambiguity affects the solution to the principal-agent problem, we reduce the model to the case of two outcomes only. In particular, we show how the properties of the optimal contract differ from the case of no ambiguity aversion, and how changes in ambiguity as well as ambiguity aversion affect the principal's profits. An additional advantage of the two-outcome case is that it is relatively easy to give a definition of the relative degree of ambiguity of two actions, which will turn out to be an important factor in determining the optimal incentive scheme.

In this section, we call outcome 2 'success', and outcome 1 'failure'. The ambiguous view of the outcome distribution for each action a can now be described by a cumulative distribution function $\mu^a(p)$ (defined for $p \in [0, 1]$), representing the confidence in potential success probabilities, or alternatively by the average success probability \bar{p}^H and the cdf $\hat{\mu}^a(\hat{p})$. The latter represents the cumulative likelihood attributed to deviations about the average success rate, its support is a subset of the interval $[-\bar{p}^a, 1 - \bar{p}^a]$. With two outcomes the payment scheme can be understood as a guaranteed payment, w_1 , and a bonus $w_2 - w_1$ that is paid only in case of success. Correspondingly we define $u_F = u(w_1)$ to be the utility level corresponding to a failure, and $u_\Delta = u(w_2) - u(w_1)$ to be the "utility bonus" in case of success.

We will focus on explaining how ambiguity affects the principal's problem of implementing a given action a . The principal wishes to minimise the costs of implementation,

$$\bar{p}^a h(u_F + u_\Delta) + (1 - \bar{p}^a) h(u_F)$$

subject to the constraints that her payment scheme given by $(u_F, u_F + u_\Delta) \in [\text{img}(u)]^2$ induces outcome a and is acceptable for the agent. After straightforward manipulations these constraints become

$$u_F + \bar{p}^a u_\Delta - c^a \geq u_0 + \underbrace{\frac{1}{\alpha} \log \left(\int \exp(-\alpha \hat{p} u_\Delta) d\hat{\mu}^a \right)}_{\mathcal{F}^a(u_\Delta)} \quad (\text{IR}) \quad (1)$$

$$(\bar{p}^a - \bar{p}^{a'}) u_\Delta \geq (c^a - c^{a'}) + \underbrace{\frac{1}{\alpha} \log \left(\frac{\int \exp(-\alpha \hat{p} u_\Delta) d\hat{\mu}^a}{\int \exp(-\alpha \hat{p} u_\Delta) d\hat{\mu}^{a'}} \right)}_{\mathcal{F}^a(u_\Delta) - \mathcal{F}^{a'}(u_\Delta)} \quad \forall a' \in A. \quad (\text{IC}) \quad (2)$$

This reformulation facilitates identifying the effects that ambiguity aversion has on the constraint set. It is easy to see that the first-best outcome is not affected by ambiguity aversion: If the action was contractible, the IC-constraint becomes obsolete and the optimal incentive scheme satisfies the IR constraint with a constant wage scheme ($u_F = u_0 + c^H$ with $u_\Delta = 0$), so that all ambiguity disappears. But as the IC constraint has to be satisfied, a non-constant incentive scheme has to be used, so that ambiguity will matter. Compared to an agent without ambiguity aversion the term $\mathcal{F}^a(u_\Delta)$ describes the effect of ambiguity aversion on the individual rationality constraint, while $\mathcal{F}^a(u_\Delta) - \mathcal{F}^{a'}(u_\Delta)$ denotes the effect on the IC constraint. If action a is ambiguous, it follows immediately that for all non-constant contracts, ambiguity aversion tightens the IR constraint, while the effect on the IC constraint could go either way. We introduce the following definition to identify the cases in which an IC constraint becomes more or less restrictive.

Definition 1 Action a' is more ambiguous than a if $\hat{\mu}^{a'}$ is a mean preserving spread of $\hat{\mu}^a$.¹³

Note that our definition implies that if action a is preferred to action a' (which is more ambiguous) by an ambiguity neutral agent under any wage contract, action a will (ceteris paribus) also be preferred by an ambiguity averse agent, as stated in the following proposition. The result is very intuitive since if an action is more ambiguous, it will become less attractive to an ambiguity averse decision maker.

Proposition 3 Suppose the principal wants to implement action a (which is not the cheapest action), and ambiguity aversion raises from zero to a positive level. Then the incentive constraint associated to action a' becomes more restrictive if a is more ambiguous, and less restrictive if a is less ambiguous than a' .

Proof The result follows from (2) since if a is more ambiguous, $\mathcal{F}^a(u_\Delta) - \mathcal{F}^{a'}(u_\Delta) \geq 0$, while if a' is more ambiguous, the inequality is reversed. \square

Yet, this observation is not enough to study the effects of ambiguity aversion on the principal's profits. The fact that it becomes harder to implement the action that is optimal in the absence of ambiguity, does not necessarily mean that the profits of the principal have to decrease. It might still be the case that another action becomes

¹³ As both distributions have the same mean, this is identical to the requirement that $\hat{\mu}^a$ second order stochastically dominates $\hat{\mu}^{a'}$.

cheaper to implement, so that the principal might actually benefit when she changes the action she implements with the optimal contract. To abstract from such concerns, we turn to the case of two actions.

Two actions Suppose there are only two actions, so that $A = \{H, L\}$ with $c^H > c^L$. It follows immediately that the low-cost action is optimally implemented with a constant incentive scheme. Thus the problem is only interesting if $\bar{p}^H > \bar{p}^L$, as otherwise the principal prefers to implement the low-cost action. Another convenient feature of this simplification is that whenever we find that the principal can implement the high-cost action in some principal-agent problem, we can easily construct another problem (by raising $q_2 - q_1$) such that implementing the high-cost action is not only efficient, but also optimal for the principal. Thus we now seek to characterise the incentive scheme that optimally implements the high-cost action. (Whether it is in fact optimal to do so will only depend on the size of $q_2 - q_1$.)

In the standard model, the optimal contract pays a positive wage bonus if the agent achieves the high outcome. Wage and bonus are determined by the unique point at which the IR constraint and the (one) IC constraint bind. Also under ambiguity (at least assuming constant absolute ambiguity attitude) we know (from above) that in the optimal solution the two constraints bind. Recall that for any $(u_F, u_F + u_\Delta) \in [\text{img}(u)]^2$ these two constraints have the form

$$u_F + \bar{p}^H u_\Delta - c^H \geq u_0 + \mathcal{F}^H(u_\Delta) \quad (\text{IR}) \quad (3)$$

and

$$(\bar{p}^H - \bar{p}^L)u_\Delta \geq c^H - c^L + \mathcal{F}^H(u_\Delta) - \mathcal{F}^L(u_\Delta). \quad (\text{IC}) \quad (4)$$

If the constraint set is empty, then it is impossible to implement the high-cost action. Otherwise the high-cost action can be implemented, and the two conditions hold with equality. However, while this is a necessary condition for an optimal incentive scheme, it is not sufficient under ambiguity aversion. While under expected utility both constraints bind only at one point (corresponding to a positive bonus), under ambiguity aversion they might also bind for a negative bonus. Under the smooth ambiguity model, it is even possible that both constraints bind for multiple contracts with a bonus of the same sign.¹⁴ Which of these should the principal pick? Lemma 2 in the Appendix states that among those contracts with the same sign it will be the one corresponding to the lowest absolute value of u_Δ .¹⁵ This is very intuitive. A lower utility bonus makes the incentive scheme both less risky and less ambiguous, and thus it becomes cheaper for the principal to satisfy the agent's IR condition. It is however not always the case that the optimal contract is one with a positive bonus.


¹⁴ This is not the case for CEU or maxmin preferences, as indifference curves are linear on each side of the 45 degree line.

¹⁵ Note that there always is such a smallest contract as the constraint set is closed due to the continuity of the utility function.

We introduce now an example that illustrates that it is possible that the IC constraint binds for both a positive and a negative value of u_Δ , and that it can be optimal to choose the negative bonus.

To understand the intuition for the optimality of the negative bonus in this example, consider first a monotonic contract (CM) corresponding to an expected utility level of u_0 (given the distribution \bar{p}^H) and an utility difference of u_Δ . If, as in this example, $\bar{p}^H < 1/2$, the potential gain $(1 - \bar{p}^H)u_\Delta$ is larger, but less likely, than the potential loss of $\bar{p}^H u_\Delta$ (high upside risk). Now compare this with an analogous non-monotonic contract (CN) with the same expected utility (under \bar{p}^H) as CM, where however the difference in utility corresponds to $-u_\Delta$. The potential gain of $\bar{p}^H u_\Delta$ realizes now with a probability of $1 - \bar{p}^H$, so it is smaller, but more likely, than the possible loss (high downside risk). Whether the monotonic or the non-monotonic contract is now cheaper for the principal depends on the agent's utility function. In particular, the non-monotonic contract can be shown to be preferred if $h''' > 0$.

Fig. 2 Negative bonus is optimal



in this example, given a monotonic contract like CM, the ambiguity averse decision maker effectively acts as if she puts a probability of somewhere between 0 and .19 (depending on the level of ambiguity aversion) to the low-cost action. Given however a non-monotonic contract like CN, she acts as if she puts a probability somewhere between .19 and .38 to the low cost action. If ambiguity aversion is high enough so that the effective probability is above .22, incentives for H can be provided with such a contract. Hence, if CM is actually the best monotonic contract, CN will probably fail to provide incentives, but a very similar contract with slightly more powerful incentives would succeed. Provided the shape of u favors the second type of contract (with high downside risk) it could be the case that even the required modification of CN remains cheaper than CM.

A similar example is discussed in Ghirardato (1994, example 1 in Sect. 3.4) for the case of the non-additive utility model, which is also driven by the fact that a less ambiguous action is implemented. We will now identify conditions that ensure that the optimal incentive scheme uses a positive bonus.

The following proposition will provide necessary conditions that ensure that the optimal contract is monotonic even under ambiguity aversion. To do so, we introduce a definition of symmetric ambiguity.

Definition 2 Ambiguity about action a is *symmetric* if the distribution $\hat{\mu}^a$ is symmetric (i.e., if $1 - \hat{\mu}^a(\hat{p}) = \hat{\mu}^a(-\hat{p})$ for all $\hat{p} \in \text{supp}(\hat{\mu}^a)$).

In words, this definition requires that the likelihood that the agent assigns to the possibility that the success probability deviates from the average success probability by more than any non-negative number \hat{p} is the same as the likelihood she assigns to a deviation of less than $-\hat{p}$.

Proposition 4 *Suppose that there are only two actions and two outcomes. Then, the optimal incentive scheme to implement the high-cost action will use a positive bonus if at least one of these three conditions holds:*

1. *The high-cost action is at least as ambiguous as the low-cost action.*
2. *It is unambiguous that the success probability of the low-cost action is lower than the average success rate for the high-cost action, in the sense that $p \in \text{supp}(\mu^L)$ implies $\bar{p}^H > p$.*
3. *Ambiguity is symmetric for both actions and*
 - *either $\bar{p}^H \geq 1/2$ and $h''' \geq 0$*
 - *or $\bar{p}^H \leq 1/2$ and $h''' \leq 0$.*

A complete proof is given in the Appendix. It proceeds by showing that, provided the low (cost) action is more ambiguous, it can be possible to implement the action with strictly higher costs even if the average success rates are the same. As discussed for the example, the sign of h''' characterizes whether or not the positive bonus is optimal in this case. The proof concludes with the observation that any increase in the average success rate of the high-cost action *ceteris paribus* favors the monotone contract.

Note that the second point implies that, if the likelihood ratio for the high-cost action is unambiguously higher than for the low-cost action, the optimal contract

 Springer

as the high-cost action is not ambiguous. U_L is the equivalent for the low-cost action. Any optimal incentive scheme corresponds to an intersection of these two curves, as has been argued in the context of Fig. 2. Finally, the curve Π (the shape of which is determined by h) depicts all incentive schemes which results in the same profits as V (assuming that H is chosen), and points below Π increase profits.

The curves U'_L and U'_H represent the same indifference curves after an increase in ambiguity aversion. As H is not ambiguous, only U'_L changes; its curvature increases. We can see that the optimal incentive scheme is given by V' in this case, and as V' lies below Π , V' is cheaper for the principal than V . Thus, the principal benefits from ambiguity aversion in this example.

The idea behind this example is intuitive: If the action we want to implement is not ambiguous, then an increase in ambiguity aversion makes it less attractive for the agent to pick the other action. Therefore, the principal can use a less high-powered incentive scheme, which is cheaper for her, due to the agent's risk aversion.

So when will it be true that an increase in ambiguity aversion leads to a decrease in profits? To keep things simple, we use the case of no ambiguity aversion as benchmark and identify in which cases ambiguity aversion is beneficial to the principal. Note that ambiguity aversion affects the set of feasible incentive schemes in two different ways: Whenever the action that is supposed to be implemented is ambiguous, and a non-constant incentive scheme is used, ambiguity aversion makes the individual rationality constraint harder to satisfy. On the other hand ambiguity aversion may have a positive or a negative effect on the incentive constraint: If the high-cost action is less ambiguous, then ambiguity aversion makes the high-cost action relatively more attractive than the low-cost action. This intuition leads to the following proposition:

Proposition 5 *Consider a principal-agent problem (P) with an ambiguity neutral agent, and the same problem but with an ambiguity averse agent (P'). Suppose that the high-cost action is more ambiguous than the low-cost action. Then the profits in problem P' can be no higher than the profits in problem P . If, instead, the high-cost action is less ambiguous, then profits may be higher in problem P' .*

Proof Assume that the high-cost action is more ambiguous. We will proof the proposition by arguing that the set of incentive schemes that implements the high-cost action in problem P' is contained in the feasible set of problem P [compare (3) and (4)]. Consider first any incentive scheme (given by u_F and u_Δ) that does not satisfy the individual rationality constraint in problem P . The IR constraint of problem P' differs by $\mathcal{F}^H(u_\Delta)$ which is non-negative. Thus, the IR constraint cannot hold in problem P' . Similarly, suppose the IC constraint was violated in P . The incentive constraint of problem P' differs by $\mathcal{F}^H(u_\Delta) - \mathcal{F}^L(u_\Delta)$. As the high-cost action is more ambiguous, $\mathcal{F}^H(u_\Delta) - \mathcal{F}^L(u_\Delta) \geq 0$, and thus the incentive constraint cannot be satisfied in problem P' as well, establishing the claim. For the case that the high-cost action is less ambiguous, assume, for example, that it is not ambiguous at all (while the low-cost action is ambiguous). Consider the optimal contract in (P). Given strict risk aversion, the IC constraint is binding. For this contract, in problem (P'), $\mathcal{F}^H(u_\Delta) = 0$ and $\mathcal{F}^L(u_\Delta) > 0$. Hence, the IR constraint holds with equality, while the IC constraint holds with a strict inequality. As the proof of Proposition 2 demonstrates, this means that a strictly better contract is available in problem (P'). \square

In principle we could try to obtain a similar statement for an increase in ambiguity aversion that starts from an already positive level. However this would require stronger assumptions about the dominance-relationship between the two actions. (The problem we face is essentially caused by the fact that CARA preferences are ordered by Arrow-Pratt risk aversion but not by the stronger relationship of Ross risk aversion. The same issue arises for CAAA preferences.) In Example 8 of the Appendix we present an example where, in the context of Proposition 5, the initial effect of introducing ambiguity aversion goes in the opposite direction of a further increase in ambiguity aversion. In this example, the high-cost action is very ambiguous compared to the low-cost action, such that the (negative) effects of ambiguity about the high-cost action are strong already for a small level of ambiguity aversion, while the positive effects of ambiguity about the low-cost action become relevant only after a further increase in ambiguity aversion.


In conclusion, we can observe that while in the standard model with two outcomes and CARA preferences an increase in absolute risk aversion reduces profits, in our model with CAAA preferences the introduction of ambiguity aversion (or a marginal increase of ambiguity aversion) may also benefit the principal. As a change in ambiguity aversion does not alter the first-best outcome, the change in the principal's problem reflects the change in the severity of the inefficiency due to moral hazard. To understand this difference between risk aversion and ambiguity aversion intuitively, note that in the standard model risk aversion is the only source of inefficiencies. Whenever the high-cost action is not very risky, but the low cost action is, the incentive problem is not very severe in the first place. For CARA preferences the positive effect of risk aversion, providing stronger incentives in a given contract, is dominated by the negative effect (a higher risk premium is required keeping incentives fixed).

4.3 Increases in ambiguity

It is easy to see that an increase in ambiguity that affects only the high-cost action can never benefit the principal (both constraints become more restrictive), while an increase in ambiguity of the low-cost action can never harm the principal (the IC becomes less restrictive). However it is less immediate whether an increase in ambiguity that affects both actions in the same way is beneficial to the principal. One might expect that such a uniform increase in ambiguity might have similar effects to an increase in ambiguity aversion. However, this turns out not to be true: While Proposition 5 showed that introducing ambiguity aversion cannot benefit the principal if the high-cost action is more ambiguous, a uniform increase in ambiguity can never benefit the principal in the opposite case, if initially the high-cost action is less ambiguous. To state this result more precisely, we introduce the following definition.

Definition 3 Suppose that for both actions $a \in \{H, L\}$, $\hat{\mu}'_a$ represents ambiguity for the principal-agent problem P' , while $\hat{\mu}_a$ describes ambiguity in P , and that the two problems are identical otherwise. Then problem P' differs from P by a uniform increase in ambiguity if ambiguity about all actions increases by the same mean preserving spread, i.e. $\hat{\mu}'_H - \hat{\mu}_H = \hat{\mu}'_L - \hat{\mu}_L = \tilde{\mu}$, where $\tilde{\mu}$ is a mean preserving spread).

Fig. 4 A uniform increase in ambiguity aversion might increase profits even if the more ambiguous action is implemented



Proposition 6 *Suppose the high-cost action is more ambiguous. Then a uniform increase in ambiguity may increase the principal's profit. However, if the high-cost action is less ambiguous, then a uniform increase in ambiguity can never be beneficial to the principal.*

The proof for the case that the high-cost action is less ambiguous is given in the Appendix. The intuition behind the result is the following: When the less ambiguous action is to be implemented, and ambiguity increases uniformly, the positive effect of ambiguity aversion on the incentive constraint remains. However, the relative difference in the ambiguity of the two actions decreases, reducing this positive effect (operating via the IC constraint). In addition, the increase in ambiguity reinforces the negative effect (via the IR constraint). On the other hand, if the more ambiguous action is to be implemented, the fact that the relative difference in ambiguity between the two action decreases, reduces the negative effect of ambiguity aversion on the IC constraint. Still, the negative effect on the IR constraint remains, so that the outcome depends on the relative size of these two effects. The example below shows that the net effect can still be positive.

Example 4 In this example, the low-cost action is not ambiguous, but the high-cost action is. V denotes the optimal incentive scheme. The primed counterparts correspond to a modified problem, which differs only by a uniform increase in ambiguity.¹⁷ In Fig. 4, the indifference curves U_H and U_L are defined as in previous examples, and V denotes the optimal incentive schemes. As Π denotes the set of all incentive schemes

¹⁷ Average success rates for H are given by $\bar{p}^H = \bar{p}'^H = .7$. The distributions $\hat{\mu}^H$ and $\hat{\mu}'^H$ are uniform distributions with support $\text{supp}(\hat{\mu}^H) = \{-.3, 0, .3\}$ and respectively, $\text{supp}(\hat{\mu}'^H) = \{-.3, .3\}$. Average success rates for L are given by $\bar{p}^L = \bar{p}'^L = .4$. In the unprimed case, there is no ambiguity about the success rate, in the primed case the deviations in $\text{supp}(\hat{\mu}'^L) = \{-.3, 0, .3\}$ are attributed with likelihoods $\mu_1^H = \mu_3^H = 1/6$ and $\mu_2^H = 2/3$. Risk attitude is given by $h(u) = u^{1.5}$.

that are equally expensive as V , and V' is below Π , the principal prefers the uniform increase in ambiguity.

Comparing Propositions 5 and 6 shows that the effects of ambiguity aversion and a uniform increase in ambiguity can be quite different. Consider for instance the case where the high-cost action is not ambiguous. In this case an increase in ambiguity aversion can never decrease the principal's payoff (even if ambiguity starts at some already positive level), but a uniform increase in ambiguity can never increase the payoff. So in this case the two changes lead to an opposite effect profits. The opposing effects of ambiguity and ambiguity aversion illustrate an advantage of working with a model that can separate ambiguity and ambiguity aversion, like the smooth ambiguity model.

5 Extensions

We now discuss the robustness of our findings to more general formulations of the principal agent problem. We begin with dropping the restriction to two outcomes, and revisit the issues of monotonicity and the effects of ambiguity and ambiguity aversion. Finally, we look at the case of multiple actions.

5.1 Multiple outcomes: monotonicity

Recall that in the absence of ambiguity the monotone likelihood ratio property (MLRP) is a sufficient condition to guarantee the monotonicity of the optimal incentive scheme, as long as there are only two actions. For the case of only two outcomes we have introduced a modification of this property for ambiguous outcome distributions that ensures that the optimal contract is monotonic even given ambiguity aversion. We will now demonstrate using an example that for more than two outcomes even a natural, but very demanding analogue of the MLRP property under ambiguity is not sufficient to ensure monotonicity of the reward scheme.

Definition 4 The actions H and L satisfy the strong monotone likelihood ratio property (SMLRP) if for all $p^H, p'^H \in \text{supp}(\mu^H)$ and $p^L, p'^L \in \text{supp}(\mu^L)$, $p_i^L/p_i^H > p'_{i+1}^L/p'_{i+1}^H$.

This is actually a very restrictive adoption of the MLRP property, as we require that the agent's likelihood ratios have to be monotone even if we use two different (possible) probability distributions to compute the likelihood ratio for two different outcomes. Note that in particular this assumption also implies that the likelihood ratios computed using the average probabilities (\bar{p}^H and \bar{p}^L) are monotonic. However, it does not suffice to ensure monotonicity, as the analyses of the example below will demonstrate. The example is motivated in an applied economic context, to illustrate that the optimality of non-monotonic schemes is not only a purely theoretical possibility.

Example 5 Consider the relationship between a firm (as a principal) and an agent who is hired to increase sales of the principal's company. That is, the agent's job is to do marketing for the principal's product. The outcome is the quantity of products sold, which is the only contractible measurement of the agent's performance. Assume the agent has access to an innovative marketing technology. Assume that the firm has a

good understanding of the distribution of sales in the absence of any marketing efforts (and this information can be made available to the agent). With respect to the effects of the marketing technology, suppose that there are two types of consumers. The firm knows how the marketing technology affects the sales for one type of consumer (type 1). This could be e.g. the younger consumers whose social media activities generate lots of data about them. There is however not much information about the remaining (type 2) consumers, so the probability that such a consumer buys after marketing efforts can be considered ambiguous.

In designing the optimal incentive scheme, the principal combines the unambiguous information about the type 1 consumers with the ambiguous information about the other types. If the chance of buying for a consumer is around .5, there could be reason to expect that intermediate sales are less ambiguous than deviations to either side of the average. We will illustrate this using the simple case of only two consumers, one of each type. In the absence of any marketing efforts, both consumers buy one unit of the good with probability $.5 - \gamma$. As motivated above, if the agent exerts effort, her marketing technology increases the probability that type 1 consumes the good to .5. In the absence of better information, principal and agent consider the three values $.5 + d$, $.5$ and $.5 - d$ as possible consumption probabilities for type 2 consumers (and they consider them to be equally likely). The parameter d could thus be interpreted as representing a dimension of ambiguity. Table 2 summarizes this example.


Note that if the agent exerts effort, the probabilities associated to a sales quantity of 0 and 2, but not to a quantity of 1, are ambiguous.

Obviously, there are many combinations of the parameters γ and d such that SMLRP holds. However, for all of them, given a fixed level of ambiguity aversion, the optimal contract will not be monotonic whenever risk aversion is sufficiently small. (Lemma 5 in the Appendix states this observation in a slightly more general context and gives a proof.)

The result is largely driven by the fact that if ambiguity and ambiguity aversion is high compared to risk aversion, the principal wants to reduce ambiguity by devising a wage scheme that does not vary a lot between the two ambiguous outcomes. Suppose $\bar{p}_2^H > \bar{p}_2^L$, as in this example and consider a payment scheme that is even constant between the two ambiguous outcomes (1 and 3, corresponding to 0 or 2 units sold). Since $\bar{p}_1^L + \bar{p}_3^L > \bar{p}_1^H + \bar{p}_3^H$, the payment for the unambiguous outcome 2 (1 unit) must be the highest in order to provide the necessary incentives. If risk aversion is strict, however, such a scheme will be, informally speaking, too risky, as in the absence of ambiguity aversion, the shape of the contract should reflect the average (increasing) likelihood ratios. In the presence of risk and ambiguity aversion, the optimal incentive

Table 2 Effort-dependent probabilities over outcomes (quantity sold)

Quantity	Low effort	High effort		
	p^L	$p_{.1}^H$	$p_{.2}^H$	$p_{.3}^H$
0	$(.5 + \gamma)^2$	$.5(.5 - d)$.25	$.5(.5 + d)$
1	$2(.25 - \gamma^2)$.5	.5	.5
2	$(.5 - \gamma)^2$	$.5(.5 + d)$.25	$.5(.5 - d)$

Fig. 5 Ambiguity and the optimal wage schedule

scheme will thus reward the highest outcome more than the lowest, but if risk aversion is small compared to ambiguity aversion, the result could be that outcome two (one unit sold) still is rewarded most.

To further illustrate which factors contribute to the optimality of non-monotonic incentive schemes, Fig. 5 (obtained using numerical methods) shows that high ambiguity (reflected by increases in the parameter d) plays a similar role as low risk aversion.¹⁸ We chose the case of $\gamma = .2$, with $\mu_1^H = \mu_3^H = .5$ which implies that SMLRP holds whenever $d < .28$. The optimal incentive scheme is non-monotonic as soon as ambiguity is large enough, even in situations where the likelihood ratios satisfy SMLRP.

With only one member of each type the example is arguably not very general, but note also when we have n members of each group, the probability assigned to selling exactly n units is still unambiguous. The fact that one outcome is not ambiguous at all is driven by the assumption that average success probabilities are .5 in both markets. Without this assumption, typically no outcome will be completely unambiguous, but a non-monotone incentive scheme might still be optimal.

Additionally, note that it may sometimes be possible for the agent to hide parts of the output that is produced (as she contributes to its creation), and in a non-monotonic incentive scheme it is in her interest to do so. In such cases the principal would not want to use a non-monotonic incentive scheme. If this additional constraint binds, the optimal incentive scheme would treat some outcomes identically, resulting in a coarser incentive scheme than one would expect in the absence of ambiguity.

Figure 5 also points to a further difference between moral hazard under risk and ambiguity. Within the expected utility framework, the optimal contract can be viewed as if it was the solution to an inference problem. The optimal contract looks as if the principal makes use of any signal that informs her about the unobservable action of the agent. Here, this is not the case since the distinction between the output level 1 and 2 is not used for a certain value of d in the optimal contract, even if this is certainly informative about the chosen action.¹⁹

¹⁸ To complete the picture, observe that increasing ambiguity aversion instead of ambiguity would have similar effects.

¹⁹ Lang (forthcoming) identifies a different channel which leads to the failure of this analogy, which is however restricted to so-called first-order models of ambiguity aversion not including the smooth ambiguity model.

5.2 Multiple outcomes: comparative statics

In our results on the effect of ambiguity and ambiguity aversion, the definition of the relationship “more ambiguous” was the only aspect specific to the case of only two outcomes. Hence, these findings generalise quite directly to the case of many outcomes if a more general notion of the relation ‘more ambiguous’ between two acts can be found. Recall that our definition of this relationship for the two-outcome case had the property that whenever two actions could be ranked, the disutility attributed to ambiguity alone [which we denoted by $\mathcal{F}^a(\mathbf{u})$] was (weakly) larger for the more ambiguous action irrespective of the wage schedule that is used.

In fact, one can also define a relation between two actions that has this property for the case of more than two actions. One possibility to do so is:

Definition 5 Let, for a fixed \mathbf{u} , $F_{(\mathbf{u}, \hat{\mu}^a)}$ denote the cdf attributed to the real-valued random variable constructed by $\sum_i \hat{p}_i u_i$, where \hat{p} is distributed according to $\hat{\mu}^a$. Action a is at least as ambiguous as a' if for every possible \mathbf{u} , $F_{(\mathbf{u}, \hat{\mu}^a)}$ is a mean preserving spread of $F_{(\mathbf{u}, \hat{\mu}^{a'})}$.

In fact, this definition is similar to the one introduced by [Jewitt and Mukerji \(2015\)](#) to compare ambiguous acts. To see an example of two actions which are ranked according to this criterion in a three outcome world, let $\hat{p} = (2\epsilon, -\epsilon, -\epsilon)$ and suppose that the support of μ^a and $\mu^{a'}$ contain at least the three possible probability distributions $p_1 = \bar{p} + \hat{p}$, $p_2 = \bar{p}$ and $p_3 = \bar{p} - \hat{p}$. If $\mu_1 = \mu_3 = \alpha$ and $\mu_2 = \beta - 2\alpha$, ambiguity increases as α (or ϵ) increases according to the above definition. Based on such a definition, we can state the following result:

Proposition 7 *Suppose there are only two actions. Suppose the high-cost action is at least as ambiguous as the low-cost action. Suppose problem \mathcal{P} and \mathcal{P}' are identical, except that the agent in \mathcal{P} is ambiguity neutral, while she is ambiguity averse in \mathcal{P}' . Then profits are (weakly) lower in problem \mathcal{P}' .*

This result includes the cases where only the low-cost action is unambiguous, and where the two actions are equally ambiguous. The proof does not differ in any essential way from the two outcome case, and is hence omitted. If H is at least as ambiguous according to the above definition, ambiguity aversion makes both constraints harder to satisfy.

Comparing ambiguity of two actions in this way has two limitations. On the one hand it is not necessarily easy to check whether two actions can be ranked according to this criterion (unless if actions are equally ambiguous or only one of them is). On the other hand, often this definition cannot rank two actions. As an example it is impossible to compare the action which is perceived to result with equal likelihood in the two probability distributions $\bar{p} \pm (\epsilon, 0, -\epsilon)$ with the action that may result in both $\bar{p} \pm (0, \epsilon, -\epsilon)$ with equal chances. However, under any monotone incentive scheme the first distribution is more vulnerable to ambiguity aversion than the latter. Hence, if the first action was considered more ambiguous our results would still apply if contracts are restricted to be monotonic (In fact, [Jewitt and Mukerji \(2015\)](#) provide a characterization of this relationship under the restriction to monotonic utility profiles as well). However, as our results suggest that if risk aversion is small compared to

ambiguity aversion, the optimal contract will not be monotonic in many cases. Hence, a broader definition of the relationship “more ambiguous than” might be particularly useful in cases where risk aversion is large, or agents have the possibility to hide output, so that the optimal contract is forced to be monotonic.

If the action chosen in the optimal contract is unambiguous, we can also show how marginal increases in ambiguity aversion affect the principal’s payoff, even if there are many outcomes.

Proposition 8 *Suppose there is no ambiguity about the action which is implemented in the optimal contract. Then any increase in ambiguity aversion (weakly) improves the payoffs of the principal (and thus reduces the severity of the incentive problem).*

The proof, located in the Appendix, exploits parallels between risk aversion and ambiguity aversion. It is intuitive since if an ambiguity averse agent finds the ambiguous off-equilibrium actions not attractive enough given a certain contract, a more ambiguity averse agent will find them even less attractive.

Recall our earlier observation that in many cases it seems likely that equilibrium actions are maybe very well understood in many instances of actual principal-agent problems, while the consequences of deviations, which never actually occur, are ambiguous. Hence, this proposition suggests that in many relevant cases the effects of increases in ambiguity aversion are very different from the effects of risk aversion, which are more typically attributed to lower profits.²⁰


Obviously an analogous result exists if the principals wants to implement an action that is ambiguous, while the other is not: In this case, an increase in ambiguity aversion can only harm the principal.

Also, working with more than two outcomes enables us to contrast the effects of increases in ambiguity about the outcome distribution with an increase in risk about the outcomes. In the example below, we show that the fact that the high-cost action becomes more risky does not mean that profits decrease.

Example 6 Suppose $Q = \{0, 1, 2\}$ and $\bar{p}^a = (\gamma, 1 - 2\gamma, \gamma)$ (for $\gamma \in [0, .5]$). As γ increases, the risk inherent in the outcome distribution increases. Yet, that does not mean that any payment scheme based on these three outcomes necessarily becomes worse for the agent: any payment scheme that treats the lower two outcomes similar and rewards the high output with the highest wage becomes in fact better for the agent as γ increases. If a is the high-cost action, and the low-cost action is described by, say, the distribution $(.5, .25, .25)$, values of γ that are close to $1/2$ will bring the principal close to the first-best outcome, even though the outcome distribution is most risky in this case. This follows as the high-cost action comes close to not having full support in this case, so that the principal can ‘punish’ the agent in a state that is likely only under the low-cost action.

²⁰ Note that this proposition remains obviously true in a case of many actions, provided the equilibrium action is the only unambiguous one.

Fig. 6 IC binds only for an action with higher costs



5.3 Multiple actions

In general, increasing the number of actions available to the agent does not change the problem in a fundamental way: The principal first checks what are the costs of implementing any possible action (where she now has to satisfy multiple incentive constraints), and then she compares these costs to the expected output the corresponding action generates. However, smooth ambiguity aversion introduces one important difference: it is no longer true that one of the incentive constraints that binds has to correspond to an action with lower costs. (This has to be true in the standard model and if ambiguity is modelled using capacities.) We illustrate using an example.

Example 7 In this example, there are two outcomes, and three actions, H , M and L . Here the medium cost action (represented by the indifference curve U_M) yields the highest expected benefits for the principal, followed by the high-cost action (U_H). The cheapest action (U_L) is also least beneficial to the principal. All of them are ambiguous, but note that the ambiguity about the high-cost (medium-benefit) action is such that it becomes significant only if the wage scheme differs a lot between the two outcomes (which is impossible in the non-additive model).²¹ If the principal wants to implement the medium-cost action, her profits are highest if she uses the incentive scheme V' (Fig. 6). In this scheme, the agent is indifferent between the medium and high-cost action, but prefers both to the low cost action. Note that the principal actually wants

²¹ The details of the example are the following: For action L , the average success rate $\bar{p}^L = .5$, deviations $\hat{\cdot} - .2, .2\}$ are equally likely. For M , $\bar{p}^M = .7$ and the deviations $\{- .15, .15\}$ are equally likely. For H , the average success rate is given by $\bar{p}^H = .675$, the deviations $\{- .275, 0, .275\}$ have likelihoods given by $\mu_1^H = \mu_3^H = .01$ and $\mu_2^H = .98$. The costs of each action are $c^L = 2$, $c^M = 2.5$ and $c^H = 2.57$, the reservation utility is $\phi(u_0) = -1$. Ambiguity attitude is CAAA with an ambiguity coefficient $\alpha = 6$. The agent is risk neutral.

to use this scheme rather than using scheme V to implement H when the benefits of the higher outcomes are large enough.

6 Conclusion

Given ambiguity about the outcome distributions, we find that allowing for ambiguity averse preferences has important consequences for the optimal contract. First, unless preferences display constant absolute ambiguity aversion, some rents may be left with the agent (the individual rationality constraint need not bind in the optimal contract). Second, we find that it will more often be the case that the optimal contract is non-monotonic. In settings where output can be hidden this means that in such cases the optimal contract will be coarse. This fits well with the often stated belief that the optimal contract is found to be more coarse than what the expected utility framework predicts. Third, we find that even in a simple model with binary outcomes increases in ambiguity aversion can lead to very different effects compared to increases in risk aversion. More risk aversion typically leads to lower profits, more ambiguity aversion can easily lead to higher profits, e.g. if there is only very little ambiguity about the action which is to be implemented, while ambiguity about possible deviations is large. Additionally, we demonstrate that, given ambiguity averse preferences, the principal's problem of finding the optimal incentive contract can no longer be viewed as an inference problem.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

Existence of an optimal contract

Lemma 1 *There exists an optimal incentive contract.*

Proof The proof follows GH (proof of Proposition 1), with the probability distribution over outcomes in the expected-utility model being replaced by the μ -average probability distribution, $\int p d\mu^a$. Observe that the low-cost action can always be implemented at finite cost. Hence it is sufficient to show that if it is possible to implement any action $a \in A$, then the problem of implementing this action at minimal cost has a solution. In fact, $\sum_i \int p_i d\mu^a u_i$ is bounded below in the optimal solution, since within the constraint set, $\phi \left(\sum_{i=1}^I \int p_i d\mu^a u(w_i) - c^a \right) \geq \int \phi \left(\sum_{i=1}^I p_i u(w_i) - c^a \right) d\mu^a \geq \phi(u_o)$. As argued in GH one can thus place artificial bounds on the constraint set in a way that ensures it becomes closed and bounded. Having assumed continuity of u and Φ , the result follows from Weierstrass' Theorem. \square

Proof of Proposition 1

(The proof of part 1) is established by the example in the text. For part 2, assume preferences are of the CAAA variety. Then we can rewrite the constraints as:

$$\begin{aligned} & - \int \left[\exp \left(-\alpha \sum_{i=1}^I p_i u(w_i) \right) \exp(\alpha c^a) \right] d\mu^a \geq -\exp(-\alpha u_0) \quad (\text{IR}') \\ & - \int \left[\exp \left(-\alpha \sum_{i=1}^I p_i u(w_i) \right) \exp(\alpha c^a) \right] d\mu^a \geq \\ & - \int \left[\exp \left(-\alpha \sum_{i=1}^I p_i u(w_i) \right) \exp(\alpha c^{a'}) \right] d\mu^{a'}, \forall a' \in A. \quad (\text{IC}') \end{aligned}$$

We use the same technique as GH: Suppose the constraint IR' does not bind, and IC' holds. Replace the payment scheme w with another payment scheme that reduces utilities given each outcome by the same amount ϵ , where ϵ is small enough so that the IR' constraint still holds. Such an ϵ exists since the image of u has no lower bound. Then also the IC' constraints still hold, as both sides of the inequality get multiplied by $\exp(\alpha\epsilon)$. But as u is increasing in w , this change clearly benefits the principal, as implementation costs decrease. \square

Proof of Proposition 2

Suppose that no IC binds in an optimal solution. Denote the optimal payment scheme by w^* , and note that it cannot be constant (otherwise, the agent would choose the least-cost action). Now consider, for some $\gamma \in (0, 1)$ an alternative payment scheme w' defined by $u(w'_i) = \gamma u(w_i^*) + (1 - \gamma) \int \left(\sum_i u(w_i^*) p_i \right) d\mu^a$. This payment scheme makes the principal strictly better off, since given strict risk aversion,

$$\begin{aligned} & \int \left(\sum_i w'_i p_i \right) d\mu^a \\ &= \int \left(\sum_i h \left(\gamma u(w_i^*) + (1 - \gamma) \int \left(\sum_{i'} u(w_{i'}^*) p_{i'} \right) d\mu^a(p') \right) p_i \right) d\mu^a(p) \\ &< \int \left(\sum_i \left(\gamma h(u(w_i^*)) p_i + (1 - \gamma) \int \left(\sum_{i'} h(u(w_{i'}^*)) p_{i'} \right) d\mu^a \right) d\mu^a \right) \\ &= \int \left(\sum_i w_i^* p_i^a \right) d\mu^a. \end{aligned}$$

Since the incentive scheme is not constant, the inequality sign follows from the strict convexity of h , which is implied by the agent's strict risk aversion,

Also, the contract w' is (weakly) beneficial to the agent, so that the individual rationality constraint continues to hold:

$$\begin{aligned} & \int \left[\phi \left(\sum_{i=1}^I \left[\gamma u(w_i^*) + (1-\gamma) \int \sum_{i'} u(w_i^*) p_{i'}' d\mu^a(p') \right] p_i^a \right) \right] d\mu^a(p) \\ &= \int \left[\phi \left(\left[\gamma \sum_{i=1}^I u(w_i^*) p_i^a + (1-\gamma) \int \sum_{i'} u(w_i^*) p_{i'}' d\mu^a(p') \right] \right) \right] d\mu^a(p) \\ &\geq \int \left[\gamma \phi \left(\sum_{i=1}^I u(w_i^*) p_i^a \right) + (1-\gamma) \phi \left(\int \sum_{i'} u(w_i^*) p_{i'}' d\mu^a(p') \right) \right] d\mu^a(p) \\ &= \gamma \int \phi \left(\sum_{i=1}^I u(w_i^*) p_i \right) d\mu^a + (1-\gamma) \int \phi \left(\sum_i u(w_i^*) p_i \right) d\mu^a \\ &= \int \left[\phi \left(\sum_{i=1}^I u(w_i^*) p_i \right) \right] d\mu^a. \end{aligned}$$

The inequality follows from concavity of ϕ . Thus w' makes the principal better off and ensures the IR constraint holds. Assuming γ is close enough to 1, continuity ensures that also all IC constraints continue to hold given w' . Hence, there is a contract w' that makes the principal better off and satisfies all constraints, contradicting the optimality of the solution w .

Proof of Proposition 4

We begin by stating a preliminary lemma.

Lemma 2 Suppose two incentive schemes $V = (u_F, u_F + u_\Delta)$ and $V' = (u'_F, u'_F + u'_\Delta)$ both satisfy the IR and IC constraint with equality and $u_\Delta u'_\Delta > 0$. Then the principal will weakly prefer incentive scheme V iff $|u_\Delta| \leq |u'_\Delta|$. Assuming either strict risk aversion or strict ambiguity aversion, she will strictly prefer V iff $|u_\Delta| < |u'_\Delta|$.

Proof Pick any two incentive schemes where $u_\Delta u'_\Delta > 0$ and both constraints bind. Assume wlog $|u_\Delta| < |u'_\Delta|$. As both schemes satisfy IR ,

$$\int \phi(u_F + \bar{p}^H u_\Delta + \hat{p} u_\Delta) d\hat{\mu}(\hat{p}) = \int \phi(u'_F + \bar{p}^H u'_\Delta + \hat{p} u'_\Delta) d\hat{\mu}(\hat{p}).$$

It follows from the concavity of ϕ that $u_F + \bar{p}^H u_\Delta \leq u'_F + \bar{p}^H u'_\Delta$. The principals implementation costs using V are given by:

$$\begin{aligned} & \bar{p}^H h(u'_F + u'_\Delta) + (1 - \bar{p}^H) h(u'_F) \\ &= \bar{p}^H h(u'_F + \bar{p}^H u'_\Delta + (1 - \bar{p}^H) u'_\Delta) + (1 - \bar{p}) h(u'_F + \bar{p}^H u'_\Delta - \bar{p}^H u'_\Delta). \end{aligned}$$

Using the above inequality,

$$\begin{aligned} & \bar{p}^H h(u'_F + \bar{p}^H u'_\Delta + (1 - \bar{p}^H)u'_\Delta) + (1 - \bar{p})h(u'_F + \bar{p}^H u'_\Delta - \bar{p}^H u'_\Delta) \\ & \geq \bar{p}^H h(u_F + \bar{p}^H u_\Delta + (1 - \bar{p}^H)u'_\Delta) + (1 - \bar{p}^H)h(u_F + \bar{p}^H u_\Delta - \bar{p}^H u'_\Delta). \end{aligned}$$

Finally, note that due to the convexity of h , and since $|u_\Delta| < |u'_\Delta|$,

$$\begin{aligned} & \bar{p}^H h(u_F + \bar{p}^H u_\Delta + (1 - \bar{p}^H)u'_\Delta) + (1 - \bar{p}^H)h(u_F + \bar{p}^H u_\Delta - \bar{p}^H u'_\Delta) \\ & \geq (1 - \bar{p}^H)h(u_F + \bar{p}^H u_\Delta - \bar{p}^H u_\Delta) + \bar{p}^H h(u_F + \bar{p}^H u_\Delta + (1 - \bar{p}^H)u_\Delta). \end{aligned}$$

The right-hand-side of this equation corresponds to the implementation costs for V . Combining the inequalities gives the result. Note that assuming either strict risk aversion or strict ambiguity aversion ensures that one of the weak inequalities is in fact strict as needed for the result. \square

We now turn to the proof of the three parts of the proposition.

Part 1 As $\mathcal{F}^H(u_\Delta) - \mathcal{F}^L(u_\Delta)$ is positive in this case, it is clear that the IC constraint in (4) cannot hold for any negative value of u_Δ .

Part 2 Let $p' = \sup(\text{supp}(\mu^L))$. Suppose $\bar{p}^H \geq p'$ but, by contradiction, $u_\Delta < 0$. Then the term $\mathcal{F}^H(u_\Delta) - \mathcal{F}^L(u_\Delta)$ as it appears in the IC constraint (4) is larger than $(p' - \bar{p}^L)u_\Delta$, since the numerator (of the fraction in the definition $\mathcal{F}^H(u_\Delta) - \mathcal{F}^L(u_\Delta)$) is at least 1 and the denominator at most $(\exp(-\alpha(p' - \bar{p}^L)u_\Delta))$. Therefore it is necessary for the IC constraint to hold that $(\bar{p}^H - p')u_\Delta > c^H - c^L$, which contradicts our assumptions.

Part 3 The proof of part 3 proceeds in the following steps. First we turn to the case of a principal that tries to implement the more costly action even though the reduced probabilities are the same. This case however is only interesting as a benchmark case, since if average probabilities are the same, the ambiguity neutral principal will always implement the low-cost action using a constant incentive scheme. Lemma 3 below will characterise the sign of the bonus in this case. Finally, we turn to the more interesting case that $\bar{p}^H > \bar{p}^L$ and argue that in this case it becomes cheaper to use the positive bonus (Lemma 4), so that the conditions we identify in Lemma 3 are still sufficient for the positive bonus, (and moreover necessary for the negative bonus), establishing the result.

Lemma 3 Suppose that $\bar{p}^H = \bar{p}^L = \bar{p}$, action H can be implemented and ambiguity about both actions is symmetric. If $\bar{p}^H > 1/2$, the principal will use the positive (negative) bonus if h''' is always positive (negative). The result is reversed if $\bar{p}^H < 1/2$. If $h''' = 0$, then positive and negative bonus are equally expensive.

Proof Consider any incentive scheme $V = (u_F, u_F + u_\Delta)$ that implements H . Let $u'_F = u_F + 2\bar{p}u_\Delta$ and $u'_\Delta = -u_\Delta$. It follows from Eqs. (3) and (4) that $V' = (u'_F, u'_F + u'_\Delta)$ satisfies both constraints iff V does.

To see this, observe that

$$\begin{aligned}\int \exp(-\alpha \hat{p}(-u_{\Delta})) d\mu^H(\hat{p}) &= \int \exp(-\alpha \hat{p}u_{\Delta}) d(1 - \mu^H(-\hat{p})) \\ &= \int \exp(-\alpha \hat{p}u_{\Delta}) d\mu^H(\hat{p})\end{aligned}$$

(and similarly for L .) Therefore plugging V' instead of V in the IC constraints leaves the RHS of both constraints unchanged. The LHS of the IR constraints (3) for V and V' are identical as well (by construction), and the LHS of (4) is zero in both cases.

Now let V be the incentive scheme that satisfies both constraints with the smallest possible positive bonus. It remains to check whether the expected wage payment associated to V is larger than the expected payment needed for V' .

To do so, rewrite the two incentive schemes in terms of the deviation about their mean, $\bar{u} = u_F + \bar{p}u_{\Delta}$. Then, $V = (\bar{u} - \bar{p}u_{\Delta}, \bar{u} + (1 - \bar{p})u_{\Delta})$ while $V' = (\bar{u} + \bar{p}u_{\Delta}, \bar{u} - (1 - \bar{p})u_{\Delta})$.

Thus, the principal prefers the positive bonus if

$$(1 - \bar{p})h(\bar{u} - \bar{p}u_{\Delta}) + \bar{p}h(\bar{u} + (1 - \bar{p})u_{\Delta}) - [(1 - \bar{p})h(\bar{u} + \bar{p}u_{\Delta}) + \bar{p}h(\bar{u} - (1 - \bar{p})u_{\Delta})] < 0.$$

The LHS of this inequality equals 0 if $u_{\Delta} = 0$. Thus the LHS is negative for $u_{\Delta} > 0$ if it is decreasing in u_{Δ} . The first derivative of the LHS w.r.t. u_{Δ} is actually given by

$$2\bar{p}(1 - \bar{p}) \left[\frac{1}{2}h'(\bar{u} + (1 - \bar{p})u_{\Delta}) + \frac{1}{2}h'(\bar{u} - (1 - \bar{p})u_{\Delta}) - \left(\frac{1}{2}h'(\bar{u} + \bar{p}u_{\Delta}) + \frac{1}{2}h'(\bar{u} - \bar{p}u_{\Delta}) \right) \right].$$

Now suppose $\bar{p} > 1/2$. Since $h'' > 0$ this derivative is negative (and thus the positive bonus is optimal) if, equivalently, h' is convex or $h''' > 0$. The relationship is reversed if either $\bar{p} < 1/2$ or h' concave, and if h' is linear, there is no difference between the positive and negative bonus.²² \square

Lemma 4 *Suppose it is optimal for the principal to implement the high-cost action using a positive bonus in a principal-agent problem P_E , where the reduced success probabilities are equal (or she is indifferent). Then she will strictly prefer the positive bonus in any principal-agent problem P that differs from P_E only by a decrease in \bar{p}^L .*


²² Thus, the principal prefers V iff an expected utility maximizer with convex utility function $-h$ prefers the (zero mean) lottery $(-\bar{p}u_{\Delta}, 1 - \bar{p}; (1 - \bar{p})u_{\Delta}, \bar{p})$, to the lottery $(\bar{p}u_{\Delta}, 1 - \bar{p}; -(1 - \bar{p})u_{\Delta}, \bar{p})$ at wealth level \bar{u} . Another way of proving the proposition could be to verify the fact that the sign of $\bar{p}^H - 1/2$ establishes which of the two lotteries third-order stochastic dominates the other, and then to apply the (unnumbered) Theorem in Whitmore (1970) to obtain the result.)

Proof Consider the cheapest scheme that implements H in problem P using a negative bonus. This incentive scheme implements H also in problem P_E : The individual rationality constraints are the same in both cases. Regarding the IC (Eq. 4), the LHS in problem P equals $(\bar{p}^H - \bar{p}^L)u_\Delta$ which is negative provided that $u_\Delta < 0$, while the LHS of P_E equals zero, which makes the constraint easier to satisfy. Thus the principal weakly prefers the best non-monotone incentive scheme in problem P_E to the best non-monotone incentive scheme of problem P . And in fact this preference must be strict as for the optimal incentive scheme of P_E both constraints must hold with equality. By assumption, this scheme is (weakly) worse than the best monotone incentive scheme of problem P . By similar arguments, the best monotone incentive scheme of problem P_E is also a feasible (monotone) scheme in problem P . Thus, this scheme is strictly better than any non-monotone incentive scheme, and thus the optimal incentive scheme must be monotone. \square

Example 8

Example 8 In this example, there is very little ambiguity about the low-cost action, while ambiguity about the high-cost action is large. The (relevant) details for the example are the following: Possible deviations about the average success probability are $\{-.2, .2\}$ for H and $\{-.2, 0, .2\}$ for L . For the high-cost action the agent puts the same likelihood on both deviations, $(\mu_1^H = \mu_2^H = .5)$, while for the low-cost action she considers the average success probability to be by far the most likely: $\mu_1^L = \mu_3^L = .5(1 - \mu_2^L) = .0005$. The graph in Fig. 7 shows, as usual, the indifference curves that corresponds to the outside option for both actions and for three different levels of ambiguity aversion: The straight lines U_H and U_L correspond to no ambiguity aversion. The corresponding dashed curves represents a positive level of ambiguity aversion (U'_H and U'_L), and an even higher level (U''_H and U''_L). For

Fig. 7 An additional increase in ambiguity aversion



wage schemes that are close to being constant, the introduction of ambiguity aversion affects only the high-cost action significantly, making it harder to implement the high-cost action. But if ambiguity aversion rises to more extreme levels, it has a sizeable effect on the low-cost action, while the additional effect on the high-cost action is negligible. The optimal incentive schemes, given by intersections of the respective two indifference curves, are labelled V , V' and V'' in increasing order of ambiguity aversion. Even if the agent is risk neutral (so that the principal's iso-profit curve is a straight line parallel to U^H), V'' is less expensive than V' despite the fact that the high-cost action is more ambiguous and V'' is the solution for a problem where the agent is more ambiguity averse than in the problem corresponding to V' . Hence implementation costs are non-monotonic in ambiguity aversion.

Proof of Proposition 6

Let unprimed variables correspond to the principal-agent problem before a uniform increase in ambiguity aversion, and their primed counterparts denote the problem after the change. Thus, for both $a \in \{H, L\}$ we can write $\hat{\mu}_{a'} = \hat{\mu}_a + \tilde{\mu}$. The new incentive constraint of the primed problem can now be rewritten as

$$(\bar{p}^H - \bar{p}^L)u_{\Delta} \geq (c^H - c^L) + \frac{1}{\alpha} \log \left(\underbrace{\frac{\int \exp(-\alpha \hat{p}_j u_{\Delta}) d\hat{\mu}^H + \int \exp(-\alpha \hat{p}_j u_{\Delta}) d\tilde{\mu}}{\int \exp(-\alpha \hat{p}_j u_{\Delta}) d\hat{\mu}^L + \int \exp(-\alpha \hat{p}_j u_{\Delta}) d\tilde{\mu}}}_{\mathcal{F}'^H(u_{\Delta}) - \mathcal{F}'^L(u_{\Delta})} \right).$$

The primed IC constraint differs from the unprimed IC only by the last term, which used to be $\mathcal{F}^H(u_{\Delta}) - \mathcal{F}^L(u_{\Delta}) = \frac{1}{\alpha} \log \left(\frac{\int \exp(-\alpha \hat{p}_j u_{\Delta}) d\hat{\mu}^H}{\int \exp(-\alpha \hat{p}_j u_{\Delta}) d\hat{\mu}^L} \right)$. Fix an arbitrary u_{Δ} that did not satisfy the incentive constraint. Suppose the high-cost action is less ambiguous. Then, as the fraction inside this term is less than 1, and since $\int \exp(-\alpha \hat{p}_j u_{\Delta}) d\tilde{\mu}^a > 0$, one can conclude that $\mathcal{F}'^H(u_{\Delta}) - \mathcal{F}'^L(u_{\Delta}) > \mathcal{F}^H(u_{\Delta}) - \mathcal{F}^L(u_{\Delta})$. Thus, the incentive constraint still does not hold. Finally, suppose only the IR constraint fails in the unprimed problem. Then, as ambiguity of the high-cost action has increased, the individual rationality constraint continues to fail in the unprimed problem. Therefore there is no incentive scheme that implements the high-cost action only in the primed problem, establishing the claim. For the case of a more ambiguous high-cost action see Example 4.

Formal analysis of Example 5

We proof a lemma to which the example fits as a special case.

Lemma 5 *Suppose that there are three outcomes, the agent is strictly ambiguity averse, only the high-cost action is ambiguous, outcome 2 is not ambiguous, and SMLRP holds. If $\bar{p}_2^L / \bar{p}_2^H \neq 1$, then the incentive scheme is non-monotonic as soon as the agent is sufficiently close to being risk neutral.*

Proof Suppose that $\bar{p}_2^H / \bar{p}_2^L > 1$. It is straightforward that the SMLRP assumption implies that $\bar{p}_3^H / \bar{p}_3^L > \bar{p}_2^H / \bar{p}_2^L > 1 > \frac{\bar{p}_1^H + \bar{p}_3^H}{\bar{p}_1^L + \bar{p}_3^L}$.

The principal solves

$$\min_u C(u) = \bar{p}_1^H h(u_1) + \bar{p}_2^H h(u_2) + \bar{p}_3^H h(u_3)$$

subject to

$$\bar{p}_1^H u_1 + \bar{p}_2^H u_2 + \bar{p}_3^H u_3 - \frac{1}{a} \log \left[\int \exp(-\alpha \hat{p}_3(u_3 - u_1)) d\mu^H \right] = u_0 + c^H \quad (5)$$

$$\bar{p}_1^L u_1 + \bar{p}_2^L u_2 + \bar{p}_3^L u_3 = u_0 + c^L. \quad (6)$$

Let $u_\Delta \equiv u_3 - u_1$. Substituting $u_3 = u_1 + u_\Delta$, the constraints can be re-expressed, using elementary algebra, as:

$$\begin{aligned} u_1 &= \frac{\bar{p}_2^H (c^L + u_0) - \bar{p}_2^L (c^H + u_0) - \bar{p}_2^L \mathcal{F}^H(u_\Delta) + u_\Delta (\bar{p}_3^H \bar{p}_2^L - \bar{p}_2^H \bar{p}_3^L)}{\bar{p}_2^H (\bar{p}_1^L + \bar{p}_3^L) - (\bar{p}_1^H + \bar{p}_3^H) \bar{p}_2^L}, \\ u_2 &= \frac{-(\bar{p}_1^H + \bar{p}_3^H) (c^L + u_0) + (\bar{p}_1^L + \bar{p}_3^L) (c^H + u_0) + (\bar{p}_1^L + \bar{p}_3^L) \mathcal{F}^H(u_\Delta) - u_\Delta (\bar{p}_3^H \bar{p}_1^L - \bar{p}_1^H \bar{p}_3^L)}{\bar{p}_2^H (\bar{p}_1^L + \bar{p}_3^L) - (\bar{p}_1^H + \bar{p}_3^H) \bar{p}_2^L}, \\ u_3 &= \frac{\bar{p}_2^H (c^L + u_0) - \bar{p}_2^L (c^H + u_0) - \bar{p}_2^L \mathcal{F}^H(u_\Delta) + u_\Delta (\bar{p}_2^H \bar{p}_1^L - \bar{p}_1^H \bar{p}_2^L)}{\bar{p}_2^H (\bar{p}_1^L + \bar{p}_3^L) - (\bar{p}_1^H + \bar{p}_3^H) \bar{p}_2^L}, \end{aligned}$$

where $\mathcal{F}^H(u_\Delta) = \frac{1}{\alpha} \log \left[\int \exp(-\alpha \hat{p}_3 u_\Delta) d\mu^H \right]$.

Denote by $(u_1(u_\Delta), u_2(u_\Delta), u_3(u_\Delta))$ the utility level corresponding to binding constraints as defined by the above equations. Note that since $\mathcal{F}^H(0) = 0$ and the denominator in the above three expressions is positive by SMLRP, one can conclude from elementary algebraic manipulations that $u_1(0) = u_3(0) < u_2(0)$ is implied by $c^H > c^L$. Also note that since $(\bar{p}_1^H + \mathcal{F}^H(u_\Delta), \bar{p}_2^H, \bar{p}_3^H - \mathcal{F}^H(u_\Delta)) \in \text{conv}(\text{supp}(\mu(a))^H)$, SMLRP ensures that $(\bar{p}_3^H - \mathcal{F}^H(u_\Delta)) \bar{p}_2^L > \bar{p}_2^H \bar{p}_3^L$ and thus $u_1'(u_\Delta) > 0$. By similar arguments, $u_2'(u_\Delta) < 0$ and $u_3'(u_\Delta) > 0$. It follows that if $u_\Delta < 0$, then $u_2(u_\Delta) > u_1(u_\Delta)$. Hence, the search for the optimal u_Δ can be restricted to the following four regions:

Region 1 ($u_\Delta < 0$)	Region 2 ($u_\Delta \geq 0$)	Region 3 ($u_\Delta > 0$)	Region 4 ($u_\Delta > 0$)
$u_3(u_\Delta) < u_1(u_\Delta) < u_2(u_\Delta)$	$u_1(u_\Delta) \leq u_3(u_\Delta) < u_2(u_\Delta)$	$u_1(u_\Delta) < u_2(u_\Delta) \leq u_3(u_\Delta)$	$u_2(u_\Delta) \leq u_1(u_\Delta) < u_3(u_\Delta)$

The principal's problem can be restated as minimizing the function $C(u_\Delta)$ defined by $C(u_\Delta) = \bar{p}_1^H h(u_1(u_\Delta)) + \bar{p}_2^H h(u_2(u_\Delta)) + \bar{p}_3^H h(u_3(u_\Delta))$.

We now show that this function is quasi-convex. Its derivative is computed (and slightly rearranged) below:

$$\begin{aligned} C'(u_\Delta) = & \bar{p}_3^L \left(\frac{\bar{p}_2^H}{\bar{p}_2^L} - \frac{\bar{p}_1^H + \bar{p}_3^H}{\bar{p}_1^L + \bar{p}_3^L} \right)^{-1} \\ & \times \left[\left(\frac{\bar{p}_3^H h'(u_3(u_\Delta))}{\bar{p}_3^L} - \frac{\bar{p}_2^H h'(u_2(u_\Delta))}{\bar{p}_2^L} \right) \left(\frac{\bar{p}_2^H}{\bar{p}_2^L} - \frac{\bar{p}_1^H + \bar{p}_3^H}{\bar{p}_1^L + \bar{p}_3^L} \right) \right. \\ & - \left(\frac{\bar{p}_2^H h'(u_2(u_\Delta))}{\bar{p}_2^L} - \frac{\bar{p}_1^H h'(u_1(u_\Delta)) + \bar{p}_3^H h'(u_3(u_\Delta))}{\bar{p}_1^L + \bar{p}_3^L} \right) \left(\frac{\bar{p}_3^H}{\bar{p}_3^L} - \frac{\bar{p}_2^H}{\bar{p}_2^L} \right) \\ & \left. + \left(\frac{\bar{p}_2^H h'(u_2(u_\Delta))}{\bar{p}_2^L} - \frac{\bar{p}_1^H h'(u_1(u_\Delta)) + \bar{p}_3^H h'(u_3(u_\Delta))}{\bar{p}_1^L + \bar{p}_3^L} \right) \frac{\mathcal{F}'(u_\Delta)}{\bar{p}_3^L} \right] \quad (7) \end{aligned}$$

So apart from a positive, multiplicative constant, C' essentially consists of three terms. First observe, as benchmark, that if h' was constant, the first two terms cancel out, so the third term remains. As \mathcal{F} is convex, it follows that in this case C has a unique minimum at the point where $\mathcal{F}^{\mathcal{H}'}(u_\Delta) = 0$, which is the case for $u_\Delta = 0$. To verify the convexity of \mathcal{F}^H we can write $\mathcal{F}^H(u_\Delta) = \frac{1}{\alpha} \log[\int \exp(-\alpha p_3^H(u_\Delta)) d\mu^H] + \bar{p}_3^H u_\Delta$. The result follows since $\exp(-\alpha p_3^H u_\Delta)$ is log-convex in u_Δ as long as $p_3^H \geq 0$, and any sum of log-convex functions is log-convex.

To evaluate the effects of risk aversion (h' is increasing), consider first region 1. For these values of u_Δ , $h'(u_3(u_\Delta)) < h'(u_1(u_\Delta)) < h'(u_2(u_\Delta))$. As $h'(u_3(u_\Delta))$ and $h'(u_1(u_\Delta))$ enter with positive coefficients in the first two terms in (7), their sum now becomes negative. Also the third term is negative (since $p_2^H / p_2^L > 1 > \frac{\bar{p}_1^H + \bar{p}_3^H}{\bar{p}_1^L + \bar{p}_3^L}$) and $\mathcal{F}'(u_\Delta) < 0$ for $u_\Delta < 0$.

Now consider region 2. While the sign of C' may vary, we argue that C is convex in that region: Applying the product rule, focus first on the effect of u_Δ via changes in instances of h' . Recall that for positive u_Δ , $u'_1(u_\Delta)$ and $u'_3(u_\Delta) > 0$ while $u'_2(u_\Delta) < 0$, and h is convex, so that it is easy to verify that the derivative of the first term of (7) is positive. Recalling that $(\bar{p}_1^H + \mathcal{F}^{\mathcal{H}'}(u_\Delta), \bar{p}_2^H, \bar{p}_3^H - \mathcal{F}^{\mathcal{H}'}(u_\Delta)) \in \text{conv}(\text{supp}(\mu^H))$, SMLRP implies that the derivative of the sum of the last two terms together is positive as well, if $\mathcal{F}^{\mathcal{H}'}$ remained constant. Finally, the effect via $\mathcal{F}^{\mathcal{H}'}$ has the same sign as $\mathcal{F}^{\mathcal{H}''}(u_\Delta)$, which is positive.

Regarding the final two regions, observe that the argument for the effect via the changes in instances of h' is as above. Regarding the effect via $\mathcal{F}^{\mathcal{H}'}$, it is positive as long as

$$\frac{\bar{p}_2^H h'(u_2(u_\Delta))}{\bar{p}_2^L} > \frac{\bar{p}_1^H h'(u_1(u_\Delta)) + \bar{p}_3^H h'(u_3(u_\Delta))}{\bar{p}_1^L + \bar{p}_3^L}.$$

As soon as this condition fails, the bounds on $\mathcal{F}^{\mathcal{H}'}(u_\Delta)$ makes it easy to see that then $C'(u_\Delta) > 0$. Also, if the condition fails for some u_Δ it fails also for all higher

values of u_Δ . Finally, observe that irrespective whether or not this condition still holds, in region 4, $C'(u_\Delta) > 0$.

Thus, $C(u_\Delta)$ is initially decreasing, then convex, and finally increasing. Thus it has a unique minimum, which is in either region 2 or region 3. Its local behavior at the boundary between those two regions characterizes the location of the optimal u_Δ .

Denote by \tilde{u}_Δ the value that makes the payment scheme constant over outcomes 2 and 3, such that $u_2(\tilde{u}_\Delta) = u_3(\tilde{u}_\Delta)$. Then it is a necessary and sufficient condition for the optimal incentive scheme to be monotonic that $C'(\tilde{u}_\Delta) < 0$. Note that $u_1(\tilde{u}_\Delta) < u_2(\tilde{u}_\Delta) = u_3(\tilde{u}_\Delta)$. Then the condition for monotonicity is equivalent to

$$\frac{\bar{p}_1^H \bar{p}_3^L (\frac{h'(u_1(\tilde{u}_\Delta))}{h'(u_2(\tilde{u}_\Delta))} - 1)}{(\bar{p}_1^L + \bar{p}_3^L) \bar{p}_2^H - \bar{p}_2^L (\bar{p}_3^H + \bar{p}_1^H \frac{h'(u_1(\tilde{u}_\Delta))}{h'(u_2(\tilde{u}_\Delta))})} \left(\frac{\bar{p}_3^H}{\bar{p}_3^L} - \frac{\bar{p}_2^H}{\bar{p}_2^L} \right) + \mathcal{F}'(\tilde{u}_\Delta) < 0. \quad (8)$$

Note that \tilde{u}_Δ does not depend on the agent's attitudes towards risk (as captured by h). The first term is negative, since $u_1(\tilde{u}_\Delta) < u_2(\tilde{u}_\Delta)$ and we assume SMLRP. We can interpret this term as the benefit (of reducing costs) due to reducing the riskiness of the wage scheme, by using a wage scheme that reflects the average likelihood ratio difference between second and the third outcome. Now consider the second term. By standard arguments, as long as $u_\Delta > 0$ we have that $\mathcal{F}^{\mathcal{H}'}(\tilde{u}_\Delta) > 0$. Thus the second term captures the costs of increasing ambiguity. Consequently, as the utility function is sufficiently close to representing risk neutral preferences, the second effect will dominate the first effect so that the optimal contract is in region 2 and hence not monotonic. The case where $\bar{p}_2^H / \bar{p}_2^L < 1$ is entirely analogous, with the exception that an optimal non-monotonic contract will now be in a region where $u_3 > u_1 > u_2$. \square

Proof of Proposition 8

Suppose that w is the optimal incentive scheme in a given principal-agent problem. Suppose the optimal contract implements action a . We proof the claim by showing that w remains feasible after an increase in ambiguity aversion. Note first that this increase does not affect the individual rationality constraint. As w satisfies each incentive constraint with equality, before the increase in ambiguity aversion,

$$\forall a' \in A, \int \phi \left(\sum_{i=1}^I p_i^{a'} u(w_i) \right) d\mu^{a'} \leq \phi \left(\sum_{i=1}^I \bar{p}_i^a u(w_i) - c^a + c^{a'} \right).$$

We can interpret this condition in a language borrowed from the literature dealing with risk aversion by noting that the certainty equivalent of the ambiguous utility lottery on the left hand side is below $\sum_{i=1}^I \bar{p}_i^a u(w_i) - c^a + c^{a'}$. What in the context of risk aversion is a (certain) monetary payment, corresponds to a certain utility level in the ambiguity framework. Similarly, a lottery offering monetary payments corresponds to a lottery over utility levels. Consider an increase in ambiguity aversion. We know from the definition of the term “more ambiguity averse”(analogous to more risk averse) that an agent that is more ambiguity averse will have an even lower certainty equivalent.

Hence all IC constraints still hold, after the increase in ambiguity aversion, and w remains feasible. \square

References

- Bose S, Renou L (2014) Mechanism design with ambiguous communication devices. *Econometrica* 82(5):1853–1872
- Carroll G (2015) Robustness and linear contracts. *Am Econ Rev* 105(2):63–536
- Di Tillio A, Kos N, Messner M (2014) The design of ambiguous mechanisms. *Rev Econ Stud* (forthcoming). <https://academic.oup.com/restud/article/84/1/237/2684495/The-Design-of-Ambiguous-Mechanisms>
- Ellsberg D (1961) Risk, ambiguity, and the Savage axioms. *Q J Econ* 75(4):643–669
- Epstein LG (2010) A paradox for the smooth ambiguity model of preference. *Econometrica* 78(6):2085–2099
- Ghirardato P (1994) Agency theory with non-additive uncertainty. Mimeo <http://web.econ.unito.it/gma/paolo/age.pdf>
- Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. *J Math Econ* 18(2):141–153
- Gollier C (2011) Portfolio choices and asset prices: the comparative statics of ambiguity aversion. *Rev Econ Stud* 78(4):1329–1344
- Grossman SJ, Hart OD (1983) An analysis of the principal-agent problem. *Econometrica* 51(1):7–45
- Herweg F, Muller D, Weinschenk P (2010) Binary payment schemes: moral hazard and loss aversion. *Am Econ Rev* 100(5):2451–2477
- Jewitt I, Mukerji S (2015) Ordering ambiguous acts. Working Paper, Oxford University
- Kellner C (2015) Tournaments as a response to ambiguity aversion in incentive contracts. *J Econ Theory* 159(Part A):627–655
- Klibanoff P, Marinacci M, Mukerji S (2005) A smooth model of decision making under ambiguity. *Econometrica* 73(6):1849–1892
- Lang M (2017) First-order and second-order ambiguity aversion. *Manag Sci*. doi:[10.1287/mnsc.2016.2443](https://doi.org/10.1287/mnsc.2016.2443)
- Lopomo G, Rigotti L, Shannon C (2011) Knightian uncertainty and moral hazard. *J Econ Theory* 146(3):1148–1172 (Incompleteness and Uncertainty in Economics)
- Maccheroni F, Marinacci M, Rustichini A (2006) Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6):1447–1498
- Mukerji S (1998) Ambiguity aversion and incompleteness of contractual form. *Am Econ Rev* 88(5):1207–1231
- Mukerji S (2003) Ambiguity aversion and cost-plus procurement contracts. Department of Economics, Discussion Paper Series 171
- Mukerji S, Tallon JM (2004) An overview of economic applications of David Schmeidler's models of decision making under uncertainty. In: Gilboa I (ed) *Uncertainty in economic theory: a collection of essays in honor of David Schmeidler's 65th birthday*, chap 13. Routledge, London
- Schmeidler D (1989) Subjective probability and expected utility without additivity. *Econometrica* 57(3):571–587
- Siniscalchi M (2009) Vector expected utility and attitudes toward variation. *Econometrica* 77(3):801–855
- Viero M-L (2012) Contracting in vague environments. *Am Econ J Microecon* 4(2):30–104
- Whitmore GA (1970) Third-degree stochastic dominance. *Am Econ Rev* 60(3):457–459