

Safety, in Numbers

Preliminary. Please Do Not Circulate

Marilyn Pease

Mark Whitmeyer*

September 8, 2023

Abstract

We introduce a way to compare actions in decision problems. An action is *safer* than another if the set of beliefs at which the decision-maker prefers the safer action increases in size (in the set inclusion sense) as the decision-maker becomes more risk averse. We provide a full characterization of this relation and discuss applications to robust belief elicitation, contracting, Bayesian persuasion, game theory, and investment hedging.

1 Introduction

Take a decision problem, in which a decision-maker (DM) chooses among the actions available to her to maximize her expected utility under some (subjective) belief. Some subset of the actions available to the DM are justifiable: she has some belief at which each of these actions is at least weakly optimal. Battigalli, Cerreia-Vioglio, Maccheroni, and Marinacci (2016) prove the striking result that increased risk aversion on the part of the DM enlarges the set of justifiable actions. That is, a

*MP: Kelley School of Business, Indiana University, marpease@iu.edu & MW: Arizona State University, mark.whitmeyer@gmail.com. We thank Tom Wiseman and Kun Zhang for their feedback.

justifiable action remains justifiable if the agent becomes more risk averse; and increased risk aversion may render optimal actions that had previously been strictly dominated.¹

However, just because an action remains optimal for some belief does not mean that the set of beliefs at which this action is optimal remains unchanged. This raises the following question. What are the properties of decision problems and actions therein for which increased risk aversion enlarges the set of beliefs at which some actions are optimal?

In this paper we fix an arbitrary decision problem and formulate a binary relation between actions available to the DM. One action i is **Safer** than another action j if the set of beliefs at which action i is preferred to j grows—in the set inclusion sense—as the agent becomes more risk averse. That is, safer actions become more attractive as the agent becomes more risk averse.

Our main result, Theorem 4.2, establishes necessary and sufficient conditions on the DM’s payoffs to two actions i and j for i to be safer than j . When there are just two states, our naïve guess was that this would be equivalent to a shallower slope of the payoff to the safer action in belief space. It turns out that this condition is too weak—our safer than relation implies a shallower slope but the converse is not true. Reassuringly, our relation is also such that if there exists a risk-free action—one with a state-independent payoff—it dominates every other action in the safer-than order.

It is natural to wonder whether this relation is also a partial order. It is not. When there are just two states, our relation is transitive, but it is not in general antisymmetric; and, therefore, not a partial order.² Importantly, symmetry of the decision problem is needed for symmetry of the relation, so in asymmetric, two-

¹By understanding a decision problem as a game with just a single player, this result is also shown by [Weinstein \(2016\)](#).

²When there are three or more states, the relation may not even be transitive.

state decision problems the relation is indeed a partial order. Nevertheless, we further show that when there are two states, for any decision problem, there exists an observationally equivalent decision problem in which our relation is a total order.

In deducing the safer-than relation, we begin with the two-state environment: there, beliefs are scalars and regions of optimality for beliefs intervals, so the proof requires only characterizing in which direction the point of indifference between the two actions moves as a result of the DM's increased risk aversion. To do this, we use an elementary result from convex analysis—the three-chord lemma. Unanticipatedly, it proves straightforward to extend the characterization of the relation to a general state space. We discover that the general case can be understood as a collection of two-state environments, and so the necessary and sufficient condition of the theorem is just that the two-state conditions hold for every pair in the collection.

The intuition behind this—that all we need to do is aggregate the two-state conditions—is as follows. We require the set of optimal beliefs for an action to grow, i.e., the set of beliefs at which the specified action is optimal in the initial decision problem must be a subset of the new set of beliefs after the agent becomes risk averse. These optimality sets are simply the intersections of half-spaces and the probability simplex. Moreover, the extreme points of these sets are vertices of the simplex and a collection of non-vertex points on the edges. These non-vertex points on the edges are precisely the indifference points between the actions for pairs of states in which each of the two actions is optimal in only one state. We need only make sure these points move in the “right direction” along the edge as the agent becomes more risk averse.

We discuss the usefulness of the safer-than order in several applications. First, we illustrate how our order can be used to eliminate beliefs in a robust elicitation setting. Even if the principal is unsure of the agent's utility function, we show

how the principal can nevertheless eliminate actions that are safer than the agent’s report. Second, we relate our order to hedging in an investment setting. Given an investor’s current holdings, we formulate a binary relation between assets: one asset “hedges” better than another if the set of beliefs justifying it expands as the investor becomes more risk averse.

Third, we identify another benefit to linear contracts in moral-hazard environments— incentive compatibility continues to hold as the agent becomes more risk averse. Fourth, we study robust persuasion, in which the principal has non-probabilistic uncertainty about the target agent’s level of risk aversion. Fifth, we note that a simple relabeling of objects allow us to compare actions for a (smoothly) ambiguity-averse DM. Sixth, and finally, we briefly discuss “safe” strategy profiles in coordination games.

1.1 Related Work

There are relatively few related papers. Closest to this one is perhaps [Whitmeyer \(2023\)](#), in which the latter of us studies transformations of decision problems that render information more valuable to a DM. Here, we study a particular variety of transformation—an increase in the DM’s risk aversion—and focus on its effect on the optimality of various actions.

This paper also harkens to the comparative statics literature; see, e.g., [Milgrom and Shannon \(1994\)](#), [Edlin and Shannon \(1998\)](#), and [Athey \(2002\)](#). We also vary a parameter, the DM’s risk aversion, and ask how this affects the DM’s behavior. However, we focus on comparisons between actions and make our relation quite demanding: the enlargement of the set of beliefs at which an action is preferred to another must arise for any monotone concave transformation of the DM’s utility.

Our relation is a sort of dominance relation between actions. “Risk dominance” was a tempting name for our “safer than” relation, but its pre-existing meaning in game theory made us prefer the latter. [Cheng and Börgers \(2023\)](#) study a gener-

alized notion of dominance that subsumes not only weak and strict dominance of strategies in games, but first- and second-order stochastic dominance of lotteries.

2 Model

There is an unknown state of the world Θ , assumed to be compact. Our protagonist is a decision-maker (DM) with a compact set of actions $A \subseteq \mathbb{R}$ (with $|A| \geq 2$), and a state dependent utility function $u: A \times \Theta \rightarrow \mathbb{R}_+$. We assume u is continuous in a , the DM has a subjective belief $\mu \in \Delta \equiv \Delta(\Theta)$, and she is a subjective expected-utility (EU) maximizer. We also specify that no action in A is weakly dominated: for all $a \in A$, there exists some $\mu \in \Delta$ at which a is uniquely optimal. We call the triple $\mathcal{D} := (A, \Theta, u(\cdot, \cdot))$ the agent's **Initial decision problem**.³

The agent becomes more risk averse if her utility function is instead \hat{u} where $\hat{u} = \phi \circ u$ for some monotone concave ϕ . We call the triple $\hat{\mathcal{D}} := (A, \Theta, \hat{u}(\cdot, \cdot))$ the agent's **Transformed decision problem**; and when the agent goes from u to \hat{u} we say she becomes more risk averse.

For any two actions $a, a' \in A$ ($a \neq a'$), in the initial decision problem, we define the set $P_{a,a'}(a)$ to be the subset of the probability simplex on which action a is weakly preferred to the a' ; formally,

$$P_{a,a'}(a) := \{x \in \Delta: \mathbb{E}_x u(a, \theta) \geq \mathbb{E}_x u(a', \theta)\}.$$

By assumption this set is non-empty and of full dimension in Δ . In the transformed decision problem, we define the set $\hat{P}_{a,a'}(a)$ in the analogous manner.

Definition 2.1. Given an initial decision problem, \mathcal{D} , action a is **Safer** than action a' if for any monotone concave ϕ , $P_{a,a'}(a) \subseteq \hat{P}_{a,a'}(a)$; i.e., the set of beliefs at which action a is preferred to a' increases in size as the DM becomes more risk averse.

³As u could (naturally) be affine, our results can be understood as statements about lotteries.

Equivalently, action a is safer than action a' if

$$\mathbb{E}_x u(a, \theta) \geq \mathbb{E}_x u(a', \theta) \Rightarrow \mathbb{E}_x \phi \circ u(a, \theta) \geq \mathbb{E}_x \phi \circ u(a', \theta)$$

for any monotone concave ϕ .

Let $i \geq_S j$ denote the binary relation **Action a is safer than Action a'** . The strict relation, $i >_S j$ denotes $i \geq_S j$ but $j \not\geq_S i$. Finally, for a fixed decision problem \mathcal{D} and some action $a \in A$, let $\mathcal{S}(a)$ denote the set of actions that are safer than a and $\mathcal{W}(a)$ denote the set of actions with respect to which a is safer:

$$\mathcal{S}(a) := \{a' \in A: a' \geq_S a\} \quad \text{and} \quad \mathcal{W}(a) := \{a' \in A: a \geq_S a'\}.$$

Similarly, let $\tilde{\mathcal{S}}(a)$ denote the set of actions that are strictly safer than a and $\tilde{\mathcal{W}}(a)$ denote the set of actions with respect to which a is strictly safer under the transformed utility function:

$$\tilde{\mathcal{S}}(a) := \{a' \in A: a' >_S a\} \quad \text{and} \quad \tilde{\mathcal{W}}(a) := \{a' \in A: a >_S a'\}.$$

3 Two States

When there are two states, 0 and 1, an agent's utility from taking any action i is captured by two numbers $\alpha(i) \equiv \alpha_i$ and $\beta(i) \equiv \beta_i$; the agent's payoffs in states 0 and 1, respectively. Moreover, as no action is weakly dominated, and as we can just relabel the actions, we specify without loss of generality that $\alpha(\cdot)$ is strictly decreasing and $\beta(\cdot)$ is strictly increasing. Note that this implies that if the DM is sure that the state is 0, she prefers lower actions to higher ones, and visa versa if she is sure that the state is 1.

When there are finitely many actions, $|A| = t$, this specializes to

$$\alpha_1 > \dots > \alpha_t, \quad \text{and} \quad \beta_1 < \dots < \beta_t.$$

3.1 “Safer” Actions

To understand what one action being “safer” than another action means, consider the DM’s choice between actions i and j . Let $x \in (0, 1)$ be the probability that the state is 1 at which she is indifferent between actions i and j ; or

$$(1 - x)\alpha_i + x\beta_i = (1 - x)\alpha_j + x\beta_j \Leftrightarrow x = \frac{\alpha_i - \alpha_j}{\alpha_i - \alpha_j + \beta_j - \beta_i},$$

where the DM chooses action i if her belief is less than x because she is relatively more certain that the state is 0 and therefore prefers the lower action. Analogously, under the transformed utility, the indifference belief is

$$\hat{x} = \frac{\phi(\alpha_i) - \phi(\alpha_j)}{\phi(\alpha_i) - \phi(\alpha_j) + \phi(\beta_j) - \phi(\beta_i)}.$$

Then, action i is safer than action j if it is chosen for more beliefs after the change in utility, or $x \leq \hat{x}$. After some algebra, this translates to

$$\frac{\phi(\beta_j) - \phi(\beta_i)}{\beta_j - \beta_i} \leq \frac{\phi(\alpha_i) - \phi(\alpha_j)}{\alpha_i - \alpha_j}. \quad (1)$$

Consider the intuition of this condition. The right-hand side is the slope of the concave transformation between α_i and α_j . In other words, it is the marginal benefit of choosing action i if the state is indeed 0 and analogously for the left-hand side. Then, equation (1) implies that if action i is safer than action j , then the marginal benefit of choosing action i if the state is 0 is larger than the marginal benefit of choosing action j if the state is 1, so that action i will be chosen more under the transformation.

Using equation (1), we can characterize exactly what the “safer than” binary relation, \geq_S , means when there are just two states:

Proposition 3.1. *For $i < j$, action i is safer than action j if and only if $\beta_i \geq \alpha_j$ and $\beta_j \geq \alpha_i$; and action j is safer than action i if and only if $\beta_i \leq \alpha_j$ and $\beta_j \leq \alpha_i$.*

Proof. Proof of this result, and others omitted from the main text, may be found in Appendix A. ■

Let us examine the conditions in the first half of Proposition 3.1. First recall that $\alpha_i > \alpha_j$ and $\beta_i < \beta_j$. Then, there are two possibilities if $\beta_i \geq \alpha_j$ and $\beta_j \geq \alpha_i$. Either $\alpha_j < \alpha_i \leq \beta_i < \beta_j$, or $\alpha_j \leq \beta_i \leq \alpha_i \leq \beta_j$. In either case, the interval $[\alpha_j, \alpha_i]$ is fully or partially to the left of $[\beta_i, \beta_j]$. Then, it is easy to see that under a concave transformation the slope between α_j and α_i will be larger than the slope between β_i and β_j , so that action i 's marginal utility is higher, and according to (1), it is the safer action.

One might assume that actions that yield comparatively consistent payoffs, regardless of the state, would be relatively "safe" actions. In other words, if we let the slope of the payoff to action i be $\gamma_i := \beta_i - \alpha_i$, then one might guess that an action with a more shallow slope, or smaller γ_i would be "safer". This is not always the case, however, as the following corollary shows.

Corollary 3.2. *If action i is safer than action j , $|\gamma_i| \leq |\gamma_j|$. The converse is not generally true.*

If action i is safer than action j , then as noted above, either $\alpha_j < \alpha_i \leq \beta_i < \beta_j$, or $\alpha_j \leq \beta_i \leq \alpha_i \leq \beta_j$. Then it is clear that $|\gamma_i| \leq |\gamma_j|$ in both cases.

To gain intuition for why a shallower slope of the payoff function is not sufficient for a safer action, consider the following example. Let $\alpha_i = 5$, $\beta_j = 4$, $\beta_i = 3$, and $\alpha_j = 1$. Then, $|\gamma_i| = 2 \leq 3 = |\gamma_j|$ so that action i has a shallower slope. This clearly violates the condition that $\alpha_i \leq \beta_j$ for i to be safer. In addition, consider the marginal benefit of choosing action i as before. The safer than relation implies that $[\alpha_j, \alpha_i]$ is to the left of $[\beta_i, \beta_j]$. In this example, however, $[\beta_i, \beta_j] \subset [\alpha_j, \alpha_i]$ so that we are no longer guaranteed the ordering of marginal benefits under the concave transformation.

3.2 Partial vs. Total Order

Lemma 3.3. *When there are two states, \succeq_S is transitive.*

\succeq_S is reflexive, but it may not be antisymmetric. This fails precisely when the decision problem is extremely symmetric. Consider, for instance, a binary-action problem that yields 1 util from action 1 in state 0 and from action 2 in state 1 and zero otherwise. By construction $\beta_1 = 0 = \alpha_2$ and $\beta_2 = 1 = \alpha_1$ so $1 \succeq_S 2$ and $2 \succeq_S 1$ but $1 \neq 2$, violating antisymmetry. It is clear; however, that extreme symmetry of the decision problem is needed for symmetry of the relation.

We say that a two-state decision problem is **Asymmetric** if there exist no actions i and $j > i$ such that $\alpha_i = \beta_j$ and $\beta_j = \alpha_i$.

Lemma 3.4. \succeq_S is antisymmetric if and only if the decision problem is asymmetric.

Proof. The decision problem is not asymmetric if and only if (by definition) there exist actions i and $j > i$ such that $\alpha_i = \beta_j$ and $\beta_j = \alpha_i$. This is true if and only if (Proposition 3.1) $i \succeq_S j$ and $j \succeq_S i$. ■

Finally, we claim that in asymmetric decision problems, \succeq_S , as constructed, is a partial order:

Theorem 3.5. When there are two states, in asymmetric decision problems, \succeq_S is a partial order.

Proof. This follows from Lemmas 3.3 and 3.4. ■

It is possible that the decision problem is such that the set of actions A is totally ordered according to \succeq_S . Our next result provides an intuitive sufficient condition, which we use later in our belief-elicitation application.

Proposition 3.6. If the DM's value function, V , is monotone on $[0, 1]$, A is totally ordered by \succeq_S .

A decision problem \mathcal{D} is **Observationally equivalent** to a decision problem $\hat{\mathcal{D}}$ if $u(a, \theta) = \hat{u}(a, \theta) + t(\theta)$, where $t: \theta \rightarrow \mathbb{R}$. Given any belief $x \in \Delta$, a DM's optimal sets of actions in any two observationally equivalent decision problems are the

same. Moreover, her value for information in two equivalent problems—regardless of whether it is endogenous or exogenous—is also the same.⁴

Proposition 3.7. *For any decision problem, there exists an observationally equivalent decision problem whose set of actions is totally ordered by \succeq_S*

Proof. Take a decision problem \mathcal{D} and the action that is uniquely optimal in state 0, which is well-defined due to our no-weakly-dominated-actions specification. We label that action 0, which yields utils α_0 and β_0 in states 0 and 1, respectively; and, hence, an expected payoff of $(1-x)\alpha_0 + x\beta_0$. Now consider an observationally equivalent decision problem $\hat{\mathcal{D}}$ where $\hat{u}(a, \theta) = u(a, \theta) + \theta(\alpha_0 - \beta_0)$. Action 0 now yields an expected payoff of α_0 , and the (necessarily convex) V is, therefore, monotone. Proposition 3.6 implies the result. ■

Example 3.8. Let $A = [0, 1]$ and $u(a, \theta) = 1 - (a - \theta)^2$.⁵ It is easy to see that $\alpha_a = 1 - a^2$ and $\beta_a = (2 - a)a$.

For $a < a' \leq \frac{1}{2}$ and for $\frac{1}{2} \leq a' < a$, it is easy to see that $a' \succ_S a$, as $\beta_i < \alpha_j$ for all $i \neq j \leq \frac{1}{2}$ and $\beta_i > \alpha_j$ for all $i \neq j \geq \frac{1}{2}$.

Now let $a < \frac{1}{2} < a'$. We have $a \succeq_S a'$ if and only if $1 - a^2 \geq (2 - a')a'$ and $1 - (a')^2 \geq (2 - a)a$, if and only if $a \geq 1 - a'$. Similarly, $a' \succeq_S a$ if and only if $a \leq 1 - a'$.

Accordingly, for an arbitrary $a \leq \frac{1}{2}$,

$$\mathcal{S}(a) = \{a' \in A : a \leq \min\{a', 1 - a'\}\} \quad \text{and} \quad \mathcal{W}(a) = \{a' \in A : a \geq \max\{a', 1 - a'\}\}.$$

⁴This is because the difference in value functions, $V(x) := \max_{a \in A} \mathbb{E}_x u(a, \theta)$ and $\hat{V}(x) := \max_{a \in A} \mathbb{E}_x \hat{u}(a, \theta) + t(\theta)$, is an affine function, so by Whitmeyer (2023) the value of information for the DM is the same in each.

⁵This is just the standard quadratic-loss specification modified by a constant so that the utils in each state are positive.

4 More Than Two States

Now let $\Theta = \{\theta_0, \dots, \theta_{n-1}\}$ ($n \geq 3$). We restrict attention to generic decision problems, in which the DM strictly prefers one of the two actions in each state.⁶ Given this, without loss of generality, we specify that action a is strictly optimal in states $\mathcal{A} := \{\theta_0, \dots, \theta_k\}$ ($0 \leq k < n - 1$) and action b is strictly optimal in states $\mathcal{B} := \{\theta_{k+1}, \dots, \theta_{n-1}\}$.

We define $\alpha_i := u(a, \theta_i)$ and $\beta_i := u(b, \theta_i)$. Accordingly, for all $\theta_i \in \mathcal{A}$, $\alpha_i > \beta_i$; and for all $\theta_j \in \mathcal{B}$, $\beta_j > \alpha_j$. Given this, we have

Theorem 4.1. *Action a is safer than action b if and only if for each $\theta_i \in \mathcal{A}$ and $\theta_j \in \mathcal{B}$, $\beta_j \geq \alpha_i$ and $\alpha_j \geq \beta_i$.*

Key to this result is the observation that we need only compare the safety of the actions in each pair of states for which a different action is optimal in each state.

Figures 1 and 2 illustrate the regions of optimality before and after the DM is made more risk-averse when one action dominates another and when one does not, respectively. In the figures, the blue region is the set of beliefs at which action 0 is preferred to 1 in the initial decision problem, and the red region is the set of beliefs at which 0 is preferred to 1 when the DM's utility is translated by $\phi(\cdot) = (\cdot)^{\frac{1}{t}}$. In both examples, there are three states. Here are two links to interactive versions of the figures, in which the agent's risk-aversion can be modified (slider t): [Example 1](#) (when $1 \succeq_S 2$) and [Example 2](#) (when $1 \not\succeq_S 2$ and $2 \not\succeq_S 1$).

4.1 General Compact State Space

It is nearly immediate to extend Theorem 4.1 to allow for Θ to be an arbitrary compact set. Indeed, for two actions $a, b \in A$ we define $\alpha(\theta) := u(a, \theta)$ and $\beta(\theta) :=$

⁶This assumption is innocuous, allowing us to save on notation and work while leaving the results unchanged.




Figure 1: When $1 \geq_s 2$ (Interactive Version).




Figure 2: When $1 \succ_S 2$ and $2 \succ_S 1$ (Interactive Version).

$u(b, \theta)$. As before, we define \mathcal{A} to be the set of states in which a is uniquely optimal; and, again focusing on generic decision problems, $\mathcal{B} = \Theta \setminus \mathcal{A}$ is the set of states in which b is uniquely optimal. Then,

Theorem 4.2. *Action a is safer than action b if and only if for each $\theta \in \mathcal{A}$ and $\theta' \in \mathcal{B}$, $\beta(\theta') \geq \alpha(\theta)$ and $\alpha(\theta') \geq \beta(\theta)$.*

Proof. Omitted, as it is identical to the proof of Theorem 4.1. ■

Some decision problems contain a risk-free action, i.e., one that guarantees a deterministic payoff to the DM. Formally, action a is risk-free if $u(a, \theta_i) = u(a, \theta_j)$ for all $\theta_i, \theta_j \in \Theta$. Keep in mind that our specification that no action is weakly dominated implies that there is at most one risk-free action. Risk-free actions interact with our relation in a natural way:

Proposition 4.3. *If there exists a risk-free action, a , $a \succ_S b$ for all $b \neq a$.*

Proof. Let a be a risk-free action, i.e., $\alpha_i \equiv \alpha(\theta_i) = \alpha(\theta_j) \equiv \alpha_j$ for all $\theta_i, \theta_j \in \Theta$. Consider some other action b . WLOG, we assume a is strictly optimal in all states $\theta_i \in \mathcal{A}$ and b is strictly optimal in all states $\theta_j \in \mathcal{B}$. Consider some e_{ij} . By construction we have $\alpha_j = \alpha_i > \beta_i$ and $\beta_j > \alpha_j = \alpha_i$, so Theorem 4.2 implies $a \succ_S b$. ■

4.2 Two Quadratic-Loss Examples

Let $\Theta = [0, 1]$, $A = [0, 1]$, and $u(a, \theta) = 1 - (a - \theta)^2$. Interestingly, \succeq_S does not have bite:

Proposition 4.4. *Under the given quadratic-loss specification, no two distinct actions can be ranked according to \succeq_S .*

Proof. Let $a < b$. Observe that the state in which the DM is indifferent between the two actions is $\hat{\theta} = \frac{a+b}{2}$. Now let us compare an arbitrary $\theta_1 \leq \hat{\theta}$ with $\theta_2 \geq \hat{\theta}$.

Observe that

$$\begin{aligned}\alpha_{\theta_1} = 1 - (a - \theta_1)^2 \geq \beta_{\theta_2} = 1 - (b - \theta_2)^2 &\Leftrightarrow \\ |b - \theta_2| \geq |a - \theta_1|,\end{aligned}$$

but neither direction of this inequality holds for all $\theta_i \in \mathcal{A} = \left[0, \frac{a+b}{2}\right]$ and $\theta_j \in \mathcal{B} = \left[\frac{a+b}{2}, 1\right]$. ■

However, a slight tweak to the decision problem, one that leaves the DM's behavior in the initial decision problem unaltered, yields not only a strict ranking over any two distinct actions, but also is one in which the ranking is antisymmetric and transitive; and, hence, one in which \geq_S totally orders the actions. Now let

$$u(a, \theta) = 1 - (a - \theta)^2 + \theta^2,$$

and consider an arbitrary pair of actions $a < b$. As before, the state in which the DM is indifferent between the two actions is $\hat{\theta} = \frac{a+b}{2}$. We compare an arbitrary $\theta_1 \leq \hat{\theta}$ with $\theta_2 \geq \hat{\theta}$:

$$\begin{aligned}\alpha(\theta_1) = 1 - a^2 + 2a\theta_1 \leq \beta(\theta_2) = 1 - b^2 + 2b\theta_2 &\Leftrightarrow \\ \theta_2 \geq \frac{b^2 - a^2}{2b} + \frac{a}{b}\theta_1,\end{aligned}$$

and

$$\begin{aligned}\alpha(\theta_2) = 1 - a^2 + 2a\theta_2 \geq \beta(\theta_1) = 1 - b^2 + 2b\theta_1 &\Leftrightarrow \\ \theta_2 \geq \frac{a^2 - b^2}{2a} + \frac{b}{a}\theta_1,\end{aligned}$$

which always hold. Consequently, $a \geq_S b$ if and only if $a \leq b$. As $a \leq b$ and $b \leq a$ if and only if $a = b$, \geq_S is reflexive. Moreover, \geq , and so \geq_S , are strongly connected and transitive, so \geq_S is a total order over actions.

4.3 Not a Partial Order

Although the safer than relation, \succeq_S , is valid for pairwise comparisons, an example suffices to establish that it is not, in general, transitive when there are three or more states.

Proposition 4.5. *If there are three or more states, \succeq_S may not be transitive.*

Proof. It suffices to construct a counterexample when there are three states. Let $\Theta = \{0, 1, 2\}$ and let there be three actions: $A = \{0, 1, 2\}$.

The payoffs to action 0 are α_0 in state 0, α_1 in state 1, and α_2 in state 2. For action 1 they are β_0 in state 0, β_1 in state 1, and β_2 in state 2; and for action 2 they are γ_0 in state 0, γ_1 in state 1, and γ_2 in state 2.

From Proposition 4.1, $0 \succeq_S 1$ if and only if $\alpha_1 \geq \beta_0$, $\beta_1 \geq \alpha_0$, $\alpha_2 \geq \beta_0$ and $\beta_2 \geq \alpha_0$. Likewise, $1 \succeq_S 2$ if and only if $\beta_2 \geq \gamma_1$, $\beta_2 \geq \gamma_0$, and $\gamma_2 \geq \beta_0$. Finally, $0 \succeq_S 2$ if and only if $\alpha_2 \geq \gamma_0$, $\gamma_2 \geq \alpha_0$, $\gamma_2 \geq \alpha_1$ and $\alpha_2 \geq \gamma_1$. It is easy to see that $0 \succeq_S 1$ and $1 \succeq_S 2$ the first three of these conditions. However, the last need not be satisfied—we can have $\gamma_1 > \alpha_2$. Indeed, the following specific values work:

$$\alpha_0 = 5, \beta_0 = 1, \text{ \& } \gamma_0 = 0; \quad \beta_1 = 6, \alpha_1 = 5, \text{ \& } \gamma_1 = 4; \quad \text{and} \quad \gamma_2 = 7, \beta_2 = 6, \text{ \& } \alpha_2 = 2.$$

■

5 Applications

Our relation is useful in a variety of settings. In §5.1, we discuss how our relation relates to robustness in elicitation. In §5.2, we formulate a notion of robust hedging. §5.3 identifies another benefit of linear contracts. In §5.4, we apply our relation to robust persuasion. In §5.5, we apply our findings to the smooth model of ambiguity aversion; and in §5.6 we formulate a notion of safe strategy profiles in games.

There are likely other interesting applications—especially promising is a robustness exercise in [Mussa and Rosen \(1978\)](#)’s screening problem—but we defer these to future work.

5.1 Robust Belief Elicitation

Let $n = 2$. There is a principal who wants to elicit a belief from an agent. The principal designs a strictly scoring rule, which yields a risk-neutral agent a value function that is a strictly convex function on $[0, 1]$. Suppose that the principal is unsure about the agent’s level of risk aversion. She knows the agent is risk-averse, so her utility is some strictly concave, monotone function over prizes, but not the exact details.

Proposition 5.1. *There exists a strictly proper scoring rule in which the agent’s set of reports is totally ordered according to \geq_S .*

Proof. For any $a \in [0, 1]$ the principal promises payoff $f(a, \theta) = 1 - (a - \theta)^2 + \mathbf{1}_{\theta=1}$. Evidently, the risk neutral agent’s payoff as a function of belief $x \in [0, 1]$ is

$$1 - x(a - 1)^2 - (1 - x)a^2 + x.$$

The linear term does not affect the agent’s optimal report so it is strictly optimal for her to be honest, i.e., $a^*(x) = x$, as required. Plugging this in, we see that $V(x) = x^2$, so Proposition 3.6 implies the result. ■

Recall that for some action $a^* \in A$, $\tilde{S}(a^*)$ denotes the set of actions that are strictly safer than a^* . Then,

Proposition 5.2. *If the agent reports a^* , she cannot have any belief $a \in \tilde{S}(a^*)$, no matter her utility function.*

5.2 Hedging

A DM's other holdings have distribution H_1 in state 1 and H_0 in state 0. Asset a pays w_1 in state 1 and w_0 in state 0, whereas asset a' pays v_1 in state 1 and v_0 in state 0. Suppose in the initial decision problem the DM is risk neutral, i.e., $u(\cdot) = \cdot$. Alternatively, we are just starting with lotteries over terminal wealth.

We say that **Asset a hedges risk better than asset a'** if for any monotone concave utility function u , the set of beliefs at which a is preferred by the DM to a' is a superset of the set of beliefs at which lottery a is preferred to a' by a risk-neutral DM.

We can put this problem in the language of the earlier framework, setting

$$\alpha_a = \int (w_0 + y) dH_0(y) = w_0 + \mu_0, \quad \text{and} \quad \beta_a = \int (w_1 + y) dH_1(y) = w_1 + \mu_1,$$

where $\mu_0 := \mathbb{E}_{H_0}[Y]$ and $\mu_1 := \mathbb{E}_{H_1}[Y]$; and

$$\alpha_{a'} = v_0 + \mu_0, \quad \text{and} \quad \beta_{a'} = v_1 + \mu_1.$$

Then, without loss of generality (as we could just relabel) we stipulate $w_0 > v_0$ and $v_1 > w_1$. Furthermore,

Proposition 5.3. *If $w_1 > v_0$, $v_1 > w_0$, and H_1 first order stochastically dominates H_0 , asset a hedges risk better than asset a' .*

The sufficient condition identified in this proposition pertains not only to the direct rewards of the two assets but also to their correlation with the DM's other holdings. The condition requires that not only is the direct reward of the better-hedging asset in the state where that asset is worse than the other, w_1 , better than its analog for the worse-hedging action, v_0 , but the distribution of the other holdings' return is also better. It gives more insurance, in a sense. It is straightforward, though tedious, to extend Proposition 5.3 to a choice between non-binary assets, so we stop here.

5.3 Another Benefit of Linear Contracts

Consider a moral-hazard problem in which an agent is hired to exert effort, a to produce some output $x = a + \theta$, where θ is a random variable with cdf F . The agent has a convex cost of effort, $c(a)$. The principal writes a contract $w: \mathbb{R} \rightarrow \mathbb{R}$ ($\rightarrow \mathbb{R}_+$ if there is limited liability), that specifies a payoff to the agent of $w(x)$ given output x .

Any contract, therefore, induces a decision problem in which the agent's set of actions is a compact feasible set of effort levels A and her utility is

$$u(a, \theta) = v(w(a + \theta) - c(a)),$$

for some concave, strictly monotone v .

A special variety of contract, is the linear contract, where $w(x) = sx + r$. The decision problem resulting from a linear contract has the following nice properties.

Lemma 5.4. *If the DM prefers action a to b in some state θ , she prefers a to b in all states.*

Lemma 5.5. *If a^* is a solution to the principal's moral-hazard problem, for any other action $b \neq a^*$, there exists a state in which the DM prefers a^* to b .*

Combining these two lemmas produces a third:

Lemma 5.6. *For any action $b \neq a^*$, the DM prefers a^* to b in all states.*

That is, a linear contract ensures that a^* is not only optimal at the agent's prior, but *weakly dominant*: for any θ and any $a \in A$,

$$v[s(a^* + \theta) + r - c(a^*)] \geq v[s(a + \theta) + r - c(a)].$$

Consequently, any monotone transformation preserves this inequality, meaning that increases in the agent's risk aversion (for one) do not affect the dominance of a^* . Naturally, then, we must have

$$\int \varphi \circ v[s(a^* + \theta) + r - c(a^*)] dF(\theta) \geq \int \varphi \circ v[s(a + \theta) + r - c(a)] dF(\theta),$$

for all $a \in A$, for any monotone concave transformation φ .

We say that IC is robust to increased risk aversion if any incentive compatible contract remains incentive compatible if the agent is more risk averse, i.e., is instead some $\varphi \circ v$. In sum,

Proposition 5.7. *In a linear contract, IC is robust to increased risk aversion.*

One pesky issue remains: the participation constraint. Note that if it binds, which it must absent limited liability, the agent is indifferent between a^* and her outside option. We say that a contract is robust to increased risk aversion if IC is robust to increased risk aversion and the participation constraint is still satisfied following any increase in the agent's risk aversion. Unfortunately, the outside option is a risk-free action, in the parlance of this paper, and so Proposition 4.3 implies the following impossibility result.

Proposition 5.8. *No contract (linear or otherwise) is robust to increased risk aversion.*

On the other hand, if there are limited liability constraints then IR need not bind. In this case, it is easy to see that there is a sort of mild robustness of linear contracts: as long as the transformation of the agent's utility is not too extreme, IR will remain satisfied, leaving the agent still willing to take the desired action a^* .

5.4 Robust Persuasion

Consider now the basic binary-state, binary-action persuasion game of [Kamenica and Gentzkow \(2011\)](#). The agent has two actions $\{0, 1\}$ and the principal always wants the agent to take action 1. In the original model, the agent's state-dependent utility function, u , is known, and in the interesting version of the example (the case in which the principal provides information to the agent), no action is weakly dominated, and the common prior is strictly lower than the agent's indifference point between the two actions, x .

In the basic setup, it does not matter which of the agent’s two actions, if any, are safe. However, suppose that the principal is unsure about the agent’s utility function. She knows the agent’s utility is in some set of utility functions \mathcal{U} , totally ordered by risk-aversion. We let the original u denote the minimal element of \mathcal{U} and \hat{u} its unique maximal element. u is the “least risk averse” the agent could be, and \hat{u} the most.

The principal is ambiguity averse and evaluates persuasion mechanisms as follows: she takes a max-min approach and plays a game against nature. She chooses a persuasion mechanism; simultaneously, nature chooses the agent’s utility from set \mathcal{U} , in order to minimize the principal’s expected payoff. We say that the principal targets utility function u' if her persuasion mechanism is the standard one given known agent utility u' .

We assume the agent’s decision problem is asymmetric. Then,

Proposition 5.9. *If the principal-desired action, 1, is safer than 0, the principal targets u . If 0 is safer than 1, the principal targets \hat{u} .*

5.5 Smooth Ambiguity

Our results extend in a natural way to the smooth ambiguity model of [Klibanoff, Marinacci, and Mukerji \(2005\)](#). Suppose our DM prefers action a to action b if and only if

$$\mathbb{E}_\nu \psi \left(\int u(a, \theta) d\pi(\theta) \right) \geq \mathbb{E}_\nu \psi \left(\int u(b, \theta) d\pi(\theta) \right),$$

where ψ is a monotone concave function and $\nu \in \Delta\Pi$ is a distribution over feasible probability measures $\pi \in \Pi \subseteq \Delta\Theta$. We specify that Π is compact.

Following [Klibanoff et al. \(2005\)](#), our DM becomes more ambiguity averse if the internal vN-M utility u stays unchanged and ψ transforms to $\hat{\psi} \equiv \phi \circ \psi$, where ϕ is some monotone concave function.

Understanding $\psi \left(\int u(a, \theta) d\pi(\theta) \right)$ as a concave functional $\psi: A \times \Delta\Theta \rightarrow \mathbb{R}_+$, we

define $\mathcal{A} \subset \Delta\Theta$ to be the set of priors at which a is uniquely optimal and $\mathcal{B} = \Delta\Theta \setminus \mathcal{A}$ to be the set in which b is uniquely optimal. We further define

$$\alpha(\pi) := \psi\left(\int u(a, \theta) d\pi(\theta)\right) \quad \text{and} \quad \beta(\pi) := \psi\left(\int u(b, \theta) d\pi(\theta)\right);$$

then, applying Theorem 4.2, obtain

Proposition 5.10. *Action a is safer than action b if and only if for each $\pi \in \mathcal{A}$ and $\pi' \in \mathcal{B}$, $\beta(\pi') \geq \alpha(\pi)$ and $\alpha(\pi') \geq \beta(\pi)$.*

5.6 A Different Kind of Risk Dominance in Games

As games are just decision problems with endogenous payoffs, our results may be applied to strategic settings. Consider the following two-player, two-action coordination game. If both players choose action a , they each get a payoff of α_1 ; and if both players choose action b , they each get a payoff of β_2 . Their mismatch payoffs are $u_1(a, b) = u_2(b, a) = \beta_1$ and $u_1(b, a) = u_2(a, b) = \alpha_2$. We assume that $\alpha_1 > \alpha_2$ and $\beta_2 > \beta_1$.

We say that a strategy pair (a_1, a_2) is **Safe** if the set of beliefs $(\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2$ with respect to which a_1 and a_2 are best responses increases in size (in the set-inclusion sense) as players become more risk-averse. Then, applying Proposition 3.1 yields

Proposition 5.11. *$\beta_1 \geq \alpha_2$ and $\beta_2 \geq \alpha_1$ if and only if (a, a) is safe.*

The necessary and sufficient condition in this proposition differs from the risk-dominance condition of Harsanyi and Selten (1988); which in this case reduces to (a, a) risk dominates (b, b) if $\beta_2 - \beta_1 \leq \alpha_1 - \alpha_2$.

References

Susan Athey. Monotone comparative statics under uncertainty. *The Quarterly Journal of Economics*, 117(1):187–223, 2002.

- Pierpaolo Battigalli, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci. A note on comparative ambiguity aversion and justifiability. *Econometrica*, 84(5):1903–1916, 2016.
- Xienan Cheng and Tilman Börgers. Dominance and optimality. *Mimeo*, 2023.
- Aaron S Edlin and Chris Shannon. Strict single crossing and the strict spence-mirrlees condition: a comment on monotone comparative statics. *Econometrica*, 66(6):1417–1425, 1998.
- John C Harsanyi and Reinhard Selten. A general theory of equilibrium selection in games. *MIT Press Books*, 1, 1988.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *The American Economic Review*, 101(6):2590–2615, 2011.
- Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.
- Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica: Journal of the Econometric Society*, pages 157–180, 1994.
- Michael Mussa and Sherwin Rosen. Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317, 1978.
- Robert R Phelps. *Convex functions, monotone operators and differentiability*, volume 1364. Springer, 2009.
- Jonathan Weinstein. The effect of changes in risk attitude on strategic behavior. *Econometrica*, 84(5):1881–1902, 2016.
- Mark Whitmeyer. Making information more valuable. *Mimeo*, 2023.

A Omitted Proofs

A.1 Proposition 3.1 Proof

Proof. By symmetry, it suffices to prove the first half of the result (i safer than j).

(\Rightarrow) We want to show that $\hat{x} \geq x$. By assumption $\beta_j \geq \alpha_i > \alpha_j$, so by the Three-chord lemma (Theorem 1.16 in Phelps (2009))

$$\frac{\phi(\alpha_i) - \phi(\alpha_j)}{\alpha_i - \alpha_j} \geq \frac{\phi(\beta_j) - \phi(\alpha_j)}{\beta_j - \alpha_j}. \quad (A1)$$

Likewise, $\beta_j > \beta_i \geq \alpha_j$ plus the Three-chord lemma imply

$$\frac{\phi(\beta_j) - \phi(\alpha_j)}{\beta_j - \alpha_j} \geq \frac{\phi(\beta_j) - \phi(\beta_i)}{\beta_j - \beta_i}. \quad (A2)$$

Combining Inequalities A1 and A2 completes this direction of the proof.

(\Leftarrow) Now suppose for the sake of contraposition that $\alpha_j > \beta_i$ (and recall $\alpha_i > \alpha_j$).

There are two possibilities: either $\alpha_i \leq \beta_j$, or $\alpha_i > \beta_j$.

Suppose first $\alpha_i \leq \beta_j$, so

$$\beta_j \geq \alpha_i > \alpha_j > \beta_i.$$

Let

$$\phi(y) = \min\{y, ky + c\},$$

where

$$c = \frac{\alpha_j(\alpha_j\beta_j - \alpha_i\beta_i)}{\alpha_j(\beta_i - \alpha_j) + \alpha_i(\alpha_j - 2\beta_i) + \alpha_j\beta_j} \text{ and } k = \frac{(\alpha_i - \alpha_j)(\alpha_j - \beta_i)}{\alpha_j(\beta_i - \alpha_j) + \alpha_i(\alpha_j - 2\beta_i) + \alpha_j\beta_j}.$$

It is straightforward to check that $k \in (0, 1)$ and $k\alpha_j + c = \alpha_j$, so ϕ is weakly concave, as required.

Moreover,

$$\frac{\phi(\beta_j) - \phi(\beta_i)}{\beta_j - \beta_i} > \frac{\phi(\alpha_i) - \phi(\alpha_j)}{\alpha_i - \alpha_j} \Leftrightarrow \frac{k\beta_j + c - \beta_i}{\beta_j - \beta_i} - \frac{k\alpha_i + c - \alpha_j}{\alpha_i - \alpha_j} > 0,$$

if and only if

$$(1 - k)(\alpha_j \beta_j - \alpha_i \beta_i) - c(\beta_j - \alpha_i + \alpha_j - \beta_i) > 0,$$

which also holds.

Finally, suppose $\alpha_i > \beta_j$, in which case we have $\alpha_i > \beta_j > \beta_i$ and $\alpha_i > \alpha_j > \beta_i$. By the three-chord lemma, we have

$$\frac{\phi(\beta_j) - \phi(\beta_i)}{\beta_j - \beta_i} \geq \frac{\phi(\alpha_i) - \phi(\beta_i)}{\alpha_i - \beta_i},$$

so it suffices to construct a concave monotone ϕ for which

$$\Psi := \frac{\phi(\alpha_i) - \phi(\beta_i)}{\alpha_i - \beta_i} - \frac{\phi(\alpha_i) - \phi(\alpha_j)}{\alpha_i - \alpha_j} > 0.$$

To that end, let

$$\phi(y) = \min\left\{y, \frac{y + \alpha_j}{2}\right\}.$$

Plugging this in, we have

$$\Psi = \frac{\frac{\alpha_i + \alpha_j}{2} - \beta_i}{\alpha_i - \beta_i} - \frac{\frac{\alpha_i + \alpha_j}{2} - \alpha_j}{\alpha_i - \alpha_j} = \frac{\alpha_j - \beta_i}{2(\alpha_i - \beta_i)} > 0,$$

as desired. ■

A.2 Corollary 3.2 Proof

Proof. (\Leftarrow) The following example suffices: $i < j$; $\alpha_i = 5$, $\beta_j = 4$, $\beta_i = 3$, and $\alpha_j = 1$.

Then, $\alpha_i > \beta_j$ so $i \not\geq_S j$ but

$$|\beta_i - \alpha_i| = 2 \leq 3 = |\beta_j - \alpha_j|.$$

(\Rightarrow) Let $i \geq_S j$, where WLOG $i < j$. By Proposition 3.1, this implies $\beta_j \geq \alpha_i$ and $\beta_j > \beta_i \geq \alpha_j$. Accordingly,

$$|\beta_j - \alpha_j| = \beta_j - \alpha_j > \beta_i - \alpha_i,$$

so if $\alpha_i \leq \beta_i$ we are done. Now let $\alpha_i > \beta_i$ and suppose for the sake of contradiction

$$\alpha_i - \beta_i > \beta_j - \alpha_j,$$

which holds if and only if

$$\alpha_i + \alpha_j > \beta_i + \beta_j,$$

which is false. ■

A.3 Lemma 3.3 Proof

Proof. Suppose in some \mathcal{D} $i \succeq_S j$ and $j \succeq_S k$, where without loss of generality each action is distinct and $i < j$.

First, we establish the following claim.

Claim A.1. $k > j$ or $k < i$.

Proof. Suppose for the sake of contradiction $i < k < j$. Then, $\alpha_k \geq \beta_j \geq \alpha_i > \alpha_k$, a contradiction. ■

Second, posit $k > j$. Then, we have $\beta_k > \beta_j \geq \alpha_i$ and $\beta_i \geq \alpha_j > \alpha_k$, which implies $i \succ_S k$. Third, posit $k < i$. Thus, $\alpha_i > \alpha_j \geq \beta_k$ and $\alpha_k \geq \beta_j > \beta_i$. ■

A.4 Proposition 3.6 Proof

Proof. First, observe that if V is monotone, \mathcal{D} must be asymmetric. Indeed, suppose for the sake of contradiction that there exist two actions i and $j > i$ with $\alpha_i = \beta_j$ and $\beta_i = \alpha_j$. Then, in the region in which i is optimal, V 's slope is $\beta_i - \alpha_i = \alpha_j - \alpha_i < 0$ and in the region in which j is optimal V 's slope is $\beta_j - \alpha_j = \alpha_i - \alpha_j > 0$. Accordingly, \mathcal{D} is not symmetric so by Lemma 3.4, \succeq_S is, therefore, antisymmetric.

Second, by Lemma 3.3, \succeq_S is transitive.

Third, suppose for the sake of contraposition that there exist $a \in A$ and $a' \in A$ that are incomparable. Without loss of generality $a < a'$. The payoffs to reporting

a and a' , respectively, as functions of belief x are

$$\alpha(1-x) + \beta x, \quad \text{and} \quad \alpha'(1-x) + \beta' x.$$

By incomparability, we must have $\beta' < \alpha$ or (exclusive) $\beta < \alpha'$. Without loss of generality, suppose the former, so $\alpha > \beta' > \beta \geq \alpha'$. Then, the slope of the value function at a belief where a is strictly optimal is $\beta - \alpha < 0$; and at a belief where a' is strictly optimal, it is $\beta' - \alpha' > 0$, so V is not monotone. ■

A.5 Theorem 4.1 Proof

Proof. Understanding $x_i := \mathbb{P}(\theta_i)$ for $1 \leq i \leq n-1$ and $1 - \sum_{i=1}^{n-1} x_i = \mathbb{P}(\theta_0)$, the DM prefers action a to b if and only if

$$\left(1 - \sum_{i=1}^{n-1} x_i\right)(\alpha_0 - \beta_0) + \sum_{i=1}^{n-1} x_i(\alpha_i - \beta_i) \geq 0.$$

By our genericity assumption, this hyperplane does not intersect the boundary of the simplex at a vertex. Moreover, there is a collection of $(k+1) \cdot (n-1-k)$ edges that connect vertices at which different actions are strictly preferred. Let \mathcal{E} be the set of such edges: for any $\theta_i \in \mathcal{A}$ and $\theta_j \in \mathcal{B}$, let e_{ij} be the edge between beliefs δ_{θ_i} and δ_{θ_j} . Thus,

$$\mathcal{E} := \{e_{ij} : \theta_i \in \mathcal{A} \wedge \theta_j \in \mathcal{B}\}.$$

For each e_{ij} , there exists a unique point on its relative interior at which the DM is indifferent between a and b , i.e., there exists $\lambda \in (0, 1)$ such that

$$\lambda\alpha_i + (1-\lambda)\alpha_j = \lambda\beta_i + (1-\lambda)\beta_j,$$

and the DM strictly prefers a (b) if $\lambda' < (>) \lambda$.

Let $\hat{\lambda}$ be the new such indifference weight in the transformed decision problem, i.e.,

$$\hat{\lambda}\phi(\alpha_i) + (1-\hat{\lambda})\phi(\alpha_j) = \hat{\lambda}\phi(\beta_i) + (1-\hat{\lambda})\phi(\beta_j).$$

It is clear that a being safer than b is equivalent to $\hat{\lambda} \leq \lambda$ for each such single dimensional comparison along edges e_{ij} .

This is because in the initial decision problem, the region of optimality for action a is the convex hull of the union of the set of beliefs $\{\delta_{\theta_0}, \dots, \delta_{\theta_k}\}$ and the set of $(k+1) \cdot (n-1-k)$ beliefs at which the DM is indifferent between the two actions. Accordingly, a being safer than b means that this set lies within the convex hull of the new such set in the transformed decision problem. It is obvious then, that it is necessary and sufficient merely to check along each edge, as if each indifferent belief lies within the convex hull of the new indifferent belief and the state θ_i , we have the required inclusion. Proposition 3.1 then yields the precise condition. ■

A.6 Proposition 5.2 Proof

Proof. Suppose for the sake of contradiction that the agent's belief is some $a \in \tilde{\mathcal{S}}(a^*)$. Without loss of generality $a < a^*$. By construction, a is uniquely optimal at belief a for a risk-neutral agent; likewise for a^* . Moreover, there exists some $\lambda \in (0, 1)$ such that $x = \lambda a + (1 - \lambda)a^*$ is the belief at which the risk-neutral agent is indifferent between reporting a and a^* .

Observe that $x > a$. By Proposition 3.1, \hat{x} , the indifference belief for a risk-averse agent, must be weakly greater than x . As $a < x \leq \hat{x}$, the risk-neutral agent must strictly prefer to report a rather than a^* , a contradiction. ■

A.7 Proposition 5.3 Proof

Proof. First,

Claim A.2. $\int [u(w_0 + y) - u(v_0 + y)] dH_0(y) \geq \int [u(w_0 + y) - u(v_0 + y)] dH_1(y)$.

Proof. Naturally, this is equivalent to

$$\int [u(v_0 + y) - u(w_0 + y)] dH_1(y) \geq \int [u(v_0 + y) - u(w_0 + y)] dH_0(y).$$

Observe that for any y ,

$$\frac{d}{dy} [u(v_0 + y) - u(w_0 + y)] = u'(v_0 + y) - u'(w_0 + y) \geq 0,$$

by the concavity of u plus the fact that $w_0 > v_0$. ■

Second,

Claim A.3. $\phi(z) := \int u(z + y) dH_1(y)$ is a concave function of z .

Proof. Directly

$$\begin{aligned} \phi(\lambda z_1 + (1 - \lambda) z_2) &= \int [u(\lambda z_1 + (1 - \lambda) z_2 + y)] dH_1(y) \\ &\geq \int [\lambda u(z_1 + y) + (1 - \lambda) u(z_2 + y)] dH_1(y), \end{aligned}$$

by the concavity of u plus the fact that everything is positive. ■

Finally, Proposition 3.1 plus Claims A.2 and A.3, yield

$$\begin{aligned} \frac{\int u(w_0 + y) dH_0(y) - \int u(v_0 + y) dH_0(y)}{w_0 - v_0} &\geq \frac{\phi(w_0) - \phi(v_0)}{w_0 - v_0} \\ &\geq \frac{\phi(v_1) - \phi(w_1)}{v_1 - w_1} \\ &= \frac{\int u(v_1 + y) dH_1(y) - \int u(w_1 + y) dH_1(y)}{v_1 - w_1}, \end{aligned}$$

as desired. ■

A.8 Lemma 5.4 Proof

Proof. Fix actions a and b with $a \neq b$. Let θ_a be a state in which the DM prefers action a , i.e.,

$$v[s(a + \theta_a) + r - c(a)] \geq v[s(b + \theta_a) + r - c(b)],$$

which holds if and only if

$$sa - c(a) \geq sb - c(b).$$

But then in any state θ , we have

$$v[s(a + \theta) + r - c(a)] \geq v[s(b + \theta) + r - c(b)],$$

by the monotonicity of v . ■

A.9 Lemma 5.5

Proof. As a^* solves the principal's problem, it must be incentive compatible. Therefore, it cannot be a strictly dominated action. Accordingly, for any other action $b \neq a^*$, there must be some state θ_{a^*} in which the DM weakly prefers a^* to b . ■

A.10 Lemma 5.6 Proof

Proof. By Lemma 5.5, for any $b \neq a^*$, the DM prefers a^* to b in some state. From Lemma 5.4, as the DM prefers a^* to b in one state, she prefers a^* to b in all states. ■

A.11 Proposition 5.9 Proof

Proof. Regardless of the distribution over posteriors induced by a persuasion mechanism G , nature prefers the threshold for taking action 1 to be as large as possible. When 1 is safer than 0 this means the agent is made maximally risk loving, i.e., she has utility function u . When 0 is safer than 1, she is made maximally risk averse, with utility \hat{u} . ■