Deliberate Randomization under Risk*

Marco Loseto † Aldo Lucia‡

February 18, 2023

Abstract

We consider a decision-maker (DM) with convex preferences who faces a set of risky actions and can delegate his choice to a randomization device. Under convexity, the DM's preferences admit a conservative multi-utility representation: each utility generates an evaluation for each action, and actions are ranked according to the lowest evaluation. Building on this multi-utility representation, we characterize the set of optimal actions and propose an efficiency criterion to rank them. Next, we narrow our attention to deliberate randomization for a DM with two utilities. In this case, we show that the DM never needs to select more than two actions with positive probability and study when the desire to randomize reveals information about risk attitude. Finally, we apply our results to games where each player has two actions and two utility functions. We show that incentives to randomize extend to strategic settings and derive a new class of mixed Nash equilibria that we call "strict" because players strictly prefer randomization. In general, convexity may lead to a multiplicity of mixed Nash equilibria. However, we show that when they exist, only strict equilibria are such that all the mixed actions are efficient.

JEL Classification: C70; C72; D81

KEYWORDS: Randomization; Convex preferences; Game Theory; Nash Equilibrium

*We thank Pierpaolo Battigalli, Simone Cerreia-Vioglio, David Dillenberger, Federico Echenique, Marco Ottaviani, Luciano Pomatto and Omer Tamuz for useful comments.

†Department of Economics, University of Chicago, mloseto@uchicago.edu.

‡Division of the Humanities and Social Sciences, Caltech, allucia@caltech.edu.

1 Introduction

Random choices interpreted as the outcome of deliberate randomization are the object of theoretical and experimental works that study decision-making under risk. The recent experimental effort to provide robust evidence about deliberate randomization motivates the growing attention on theoretical models that can rationalize this observed pattern.[1] Yet, while several appealing models belong to this category, the lack of analytical tractability that characterizes a large class of them limits their use in applied research.

Convexity is the axiom that captures preferences for randomization. It requires that if a decision-maker (DM) is indifferent between two lotteries p and q, then any convex combination between p and q is weakly preferred. Preferences that satisfy convexity and few other rationality requirements admit a conservative multi-utility representation: the DM has a set of utility functions and reacts to this multiplicity by evaluating each lottery with the utility function that yields the lowest payoff.[2] One can imagine either a DM with multiple selves or a Rawlsian planner that aggregates the preferences of different individuals.[3]

In many economic problems, a DM chooses an action from a set of available alternatives to maximize his well-being. Unfortunately, the conservative multi-utility representation is not differentiable, and consequently, it is not possible to use standard optimization techniques to characterize the properties of the set of optimal actions. To overcome this issue, we provide a general characterization of the DM's optimal action(s) in terms of the strict upper-contour sets of the utilities involved in the representation. In particular, we show that an action maximizes the DM's preferences if and only if the intersection of the strict upper-contour sets of the "worst-off" utilities (i.e., the ones whose evaluation of the DM's action is the

[1] See for instance Agranov and Ortoleva (2017).

[2] Cerreia-Vioglio (2009) first studies the implications of convexity for preferences under risk.

[3] Cerreia-Vioglio (2009) describes a DM with multiple selves which is unsure about one or possibly all of the following: the value of the decision outcome, future tastes, and degree of risk aversion.

lowest) is empty.[4]

Furthermore, we propose a notion of efficiency that strengthens the requirement of optimality in two ways. We start by calling an optimal action minimal if no other action constitutes a Pareto improvement for the set of worst-off utilities that it induces. We motivate this additional requirement by showing that minimal actions induce the smallest set of worst-off utilities. Next, we define an action efficient if it is minimal and there is no other action that constitutes a Pareto improvement for the set of all utilities.[5] We prove that there is always an efficient action within the set of optimal actions. Consequently, this efficiency notion can always serve as a selection criterion for the case of multiple optimal actions. Moreover, we derive conditions that guarantee the uniqueness of the optimal action.

Our general analysis of the set of optimal actions and their properties lays the groundwork for the second part of our paper, where aiming for higher tractability, we turn to the analysis of deliberate randomization for a DM with two utilities. In this setting, the DM never finds it optimal to select more than two actions with positive probability. The value of this result is twofold. First, it shows that finding optimal actions is easier in the two utility specification than in the generic finite case because it is enough to focus on randomization over at most two actions. Second, it provides a testable implication of this assumption: a DM with two utilities should never be willing to pay any positive monetary amount to pick more than two actions with positive probability.

We call randomization strictly beneficial when it allows the DM to achieve a strictly greater payoff than with any pure action. If the DM is indifferent among all the pure actions, randomization is always strictly beneficial unless preferences do

[4] Properties of maxmin optima have also been exploited in other contexts. For example, in the theory of optimal auctions, Chung and Ely (2007) provide sufficient conditions for dominant-strategy mechanisms to have maxmin foundation.

[5] In the auction context, Börgers (2017) refines Chung and Ely (2007)'s criterion in order to exclude dominant-strategy mechanisms that he classifies as "dominated". Our refinement notion is stronger because we require efficient actions to be minimal.

not degenerate to expected utility. In this scenario, we explicitly derive the support of the optimal random choices. Moreover, we study the case of two pure actions, which is relevant in experimental settings. We then use this result to characterize when an observed preference for randomization can be used to rule out both risk aversion and risk seeking attitudes in the cautious expected utility (C-EU) model, which is a special case of the class of preferences that we consider.[6]

Finally, we apply our results to non-cooperative game theory, the main analytical tool to build formal economic models. One obstacle in studying games in which players have non-expected utility preferences is that the notion of Nash equilibrium often needs to be modified. For instance, the Nash requirement of correct conjectures is not well-defined in models under uncertainty with multiple beliefs. Instead, the class of convex preferences that we focus on does not feature the same problem. While each player has multiple utility functions, the conjecture is unique. Moreover, all the standard assumptions for the existence of a Nash equilibrium hold for the class of convex preferences that we consider.

Specifically, we study a static game with two players with convex preferences. Each player has two actions and two utility functions. We partition the possible mixed equilibria into three categories: weak, partially strict, and strict. In weak equilibria, players are indifferent in equilibrium between their mixed and pure actions, as in the expected utility case. Partially strict and strict mixed equilibria instead constitute an element of novelty. In these equilibria, at least one player strictly prefers the equilibrium mixed action to the pure actions in the support. We provide necessary conditions for the existence of these new types of equilibria and we illustrate them in a simple coordination game. In this example, convexity may lead to a multiplicity of mixed Nash equilibria. However, we show that when they exist, only strict mixed Nash equilibria are such that both players play efficient mixed actions.

[6] See Cerreia-Vioglio et al. (2015).

1.1 Related literature

This paper contributes to the recent theoretical literature that studies stochastic choice as the outcome of deliberate randomization.[7] This strand of contributions builds on the idea first proposed by Machina (1985) that individuals with non-stochastic preferences over lotteries may have an explicit desire to randomize their choices. Battigalli et al. (2017) develop a framework to model random choices under uncertainty. Our paper, instead, focuses on choices under risk, building on the multi-utility representation result obtained by Cerreia-Vioglio (2009) for preferences that satisfy convexity. This representation is appealing because it encompasses several well-known decision criteria under risk, such as the cautious expected utility (C-EU) model of Cerreia-Vioglio et al. (2015) or the maxmin model of Maccheroni (2002).

The premise of this paper is that the multi-utility representation in Cerreia-Vioglio (2009) is not differentiable, so standard optimization techniques to study random choices are not viable. Cerreia-Vioglio et al. (2020) make an analogous remark for betweenness preferences that satisfy Dillenberger's (2010) negative certainty independence axiom, such as the Gul's (1991) model of disappointment aversion. Given that negative certainty independence implies convexity, the representation in Cerreia-Vioglio (2009) is more general. At the same time, our focus on a finite set of utilities in practice allows for betweenness violations. For this reason, we see our paper as complementary to Cerreia-Vioglio et al. (2020) for the analysis of preferences in which randomization can be strictly beneficial.

A growing experimental literature supports the hypothesis that subjects make stochastic decisions deliberately. Agranov and Ortoleva (2017) provide evidence in favor of the class of convex preferences that we consider, showing that models of bounded rationality or random preferences cannot rationalize subjects' stochastic behavior in their experiment. Consistently with the conservative multi-utility

[7] See Agranov and Ortoleva (2022) for a recent review of the literature.

interpretation of convex preferences, hedging and diversification were the main motivations behind this stochastic behavior. Agranov and Ortoleva (2020) push this observation further, showing not only the existence of questions for which subjects want to randomize but also their prevalence. Our paper provides new testable predictions for models of deliberate randomization, deriving properties of optimal random choices and conditions under which strict preferences for randomization are inconsistent with both risk aversion and risk seeking attitudes.

Evidence of preferences for randomization extends to strategic settings. Agranov et al. (2021) show that randomization is a stable and pervasive feature of several choice environments, including games. In their experiment, a sizable part of individuals displays preferences for randomization in individual decision problems but especially in games. Calford (2021) studies the role of mixed actions for ambiguity averse players with maxmin expected utility preferences (Gilboa and Schmeidler (1989)). He proves that the set of rationalizable strategies grows larger as preferences for randomization weaken. We also apply our results to static games. However, while each player has multiple utilities in our setting, the conjecture is unique. Consequently, it is not necessary to modify the Nash equilibrium notion, as is the case with models under ambiguity.[8]

Allen and Rehbeck (2021) also study preferences for randomization in settings of strategic interaction by focusing on concave perturbed utility games. In their framework, players' preferences are represented by a general base utility index and an additively separable concave perturbation function. By making different functional form assumptions on the perturbation function, they construct and study properties of the best response functions. Our framework differs because rather than relying on utility perturbations, we start by imposing convexity on players' preferences and exploit the general axiomatic representation in Cerreia-Vioglio (2009) to model preferences for randomization. At the same time, in Section 6,

[8]See, for instance, Marinacci (2000).

we provide a closed-form expression for the best response function under the assumption that players have maxmin preferences and that randomization is strictly beneficial. We then use this characterization to compute the set of all possible Nash equilibria in a simple coordination problem.

1.2 Outline

The rest of the paper is organized as follows. Section 2 sets up the decision model. Section 3 provides a general characterization of the set of DM's optimal actions. Section 4 deals with the efficiency and uniqueness of optimal actions. Section 5 studies the implications of deliberate randomization for a DM with two utilities. Section 6 applies our results to the analysis of mixed Nash equilibria in a static game where players have convex preferences. Section 7 summarizes the main findings and concludes. All the proofs of the statements are collected in the appendix.

2 Model

This section begins with the introduction of the decision framework. After that, we describe the conservative multi-utility model of Cerreia-Vioglio (2009) and the additional assumptions we impose on his representation.

2.1 Decision framework

Following Luce and Raiffa (1957),[9] a decision framework is a quartet $\langle A, S, C, \rho \rangle$, where A is a finite set of conceivable pure actions, S is a finite set of states, C is a finite set of consequences and $\rho\colon A \times S \to C$ is the consequence function. Given a generic set X, we denote by $\triangle(X)$ the set of probability distributions over X. The

[9] Luce and Raiffa (1957) introduce this framework to study choice under uncertainty. Here, we endow the DM with a subjective belief over the states.

DM can commit his actions to some random devices. We denote by $\mathcal{A} = \triangle(A)$ the set of feasible actions.[10]

The DM has a belief $\mu \in \triangle(S)$ over the states. Every feasible action α, given a belief μ induces a lottery according to the stochastic outcome function:

$$\hat{\rho}\colon \mathcal{A} \times \triangle(S) \to \triangle(C).$$

The specification of the belief is relevant in Section 6, where we consider an application of our results to game theory. In all the other sections, we omit the dependence from the belief in the notation because it plays no specific role.

2.2 Preferences

We denote by $\mathbb{E}(\alpha, v)$ the expected utility of action α, with utility $v\colon C \to \mathbb{R}$:

$$\mathbb{E}(\alpha, v) = \sum_{a \in A} \alpha(a) \sum_{s \in S} \mu(s) v\left(\rho(a, s)\right).$$

We also indicate by $\succsim_v$ the binary relation representing the preferences of an expected utility DM with utility v:

$$\alpha \succsim_v \beta \Leftrightarrow \mathbb{E}(\alpha, v) \geq \mathbb{E}(\beta, v),$$

with $\alpha, \beta \in \mathcal{A}$. Moreover, we denote by $\succ_v$ and $\sim_v$ the asymmetric and symmetric parts of $\succsim_v$, respectively. Given an action $\alpha \in \mathcal{A}$, we denote by $SUCS_v(\alpha)$ the strict upper contour set of α for utility v. That is,

$$SUCS_v(\alpha) := \{\alpha' \in \mathcal{A} \,:\, \alpha' \succ_v \alpha\}.$$

In words, $SUCS_v(\alpha)$ is the set of actions that utility v strictly prefers to α.

When a preference $\succsim$ over $\mathcal{A}$ is complete, transitive, continuous and satisfies convexity,[11] Cerreia-Vioglio (2009) shows the existence of a utility function u that

[10] $\triangle(A)$ is the set of conceivable actions. In principle, not all conceivable actions are feasible: $\mathcal{A} \subseteq \triangle(A)$. In this paper, we assume $\mathcal{A} = \triangle(A)$.

[11] The preference relation $\succsim$ satisfies convexity if and only if for all $\alpha \in \mathcal{A}$, $\beta \in \mathcal{A}$ and $\lambda \in (0, 1)$, $\alpha \sim \beta \Rightarrow \lambda\alpha + (1 - \lambda)\beta \succsim \alpha$.

represents $\succsim$ as follows: there exist a set of normalized utility functions $\mathcal{W}$ and a function $\mathrm{U}\colon \mathbb{R} \times \mathcal{W} \to [-\infty, +\infty]$ such that for all $\alpha \in \mathcal{A}$,[12]

$$u(\alpha) = \inf_{v \in \mathcal{W}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right]. \qquad (\star)$$

For every utility $v \in \mathcal{W}$, the DM computes the expected utility of action α and then distorts it with the function $\mathrm{U}[\cdot, v]$, which we assume to be strictly increasing in the first argument.[13] Of all possible distorted expected utility evaluations, the DM adopts a conservative criterion assigning to α the smallest one. We further assume that $\mathcal{W}$ is finite so that the smallest evaluation is always well-defined. We call an action optimal if it maximizes ($\star$).

Because this representation relies on minimal assumptions for $\succsim$, it encompasses several decision models under risk. If $\mathcal{W}$ is a singleton, the representation reduces to expected utility. When $\mathrm{U}[x, v] = x$ for all $v \in \mathcal{W}$ and $x \in \mathbb{R}$, we obtain the maxmin expected utility model of Maccheroni (2002). Finally, if $\mathrm{U}[x, v] = v^{-1}(x)$ for all $v \in \mathcal{W}$ and $x \in \mathbb{R}$, we get the cautious expected utility model of Cerreia-Vioglio et al. (2015).

Given an action $\alpha \in \mathcal{A}$, denote by S_α the support of α and by M_α the set of worst-off utilities that α induces:

$$M_\alpha := \arg\min_{v \in \mathcal{W}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right].$$

Moreover, given a utility function $v \in \mathcal{W}$, denote by P_v the set of pure actions for which v belongs to the induced set of worst-off utilities:

$$P_v := \{a \in A | v \in M_a\}.$$

[12] We fix an arbitrary consequence $c \in C$ and define $\mathcal{W} = \mathcal{W}_1 = \{v \in \mathbb{R}^C : v(c) = 1\}$.

[13] Cerreia-Vioglio (2009) proves that $\mathrm{U}[\cdot, v]$ must be increasing in the first argument. The additional requirement that we impose is satisfied in the special cases of Maccheroni (2002) and Cerreia-Vioglio et al. (2015).

3 Optimal actions

Our main result characterizes the set of optimal actions in terms of the strict upper contour sets of the worst-off utilities that these actions induce.

Proposition 1. *Action* $\alpha^* \in \mathcal{A}$ *is optimal if and only if* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha^*) = \emptyset$.

Proposition 1 establishes that an action is optimal whenever there is no other action that is strictly better for all the worst-off utilities that the action induces. Suppose that the intersection of the strict upper contour sets of all the worst-off utilities that action α^* induces is empty. Then, for all actions $\alpha \in \mathcal{A}$, there must exists a utility $v \in M_{\alpha^*}$ such that $\alpha^* \succsim_v \alpha$. Consequently, action α^* is optimal.

For the other direction, suppose that the intersection of the strict upper contour sets of all the worst-off utilities that action α^* induces is non-empty. Then, there must exist an action α that is strictly better than α^* for all utilities $v \in M_{\alpha^*}$. Given that the set of utilities $\mathcal{W}$ is finite, the payoffs of action α^* for utilities that do not belong to M_{α^*} must be larger than the payoff of the worst-off utilities in M_{α^*} by some finite amount, say $\epsilon > 0$. Because all utilities are continuous, it is possible to mix action α^* with a little bit of α to make all utilities in M_{α^*} better-off without rendering anyone outside M_{α^*} worst-off by more than ϵ. Therefore, action α^* is not optimal.

Proposition 1 hints at a strategy to verify whether an action α^* is optimal: check whether the intersection of the strict upper contour sets of all the worst-off utilities in M_{α^*} is empty. The next proposition introduces an indirect tool to simplify this task.

Proposition 2. *Action* $\alpha^* \in \mathcal{A}$ *is optimal if and only if* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha) = \emptyset$ *for all* $\alpha \in \mathcal{A}$ *with* $S_\alpha \subseteq S_{\alpha^*}$.

According to Proposition 2, an action α^* is optimal whenever there are no actions α and $\tilde{\alpha}$ such that the support of action α^* contains the support of action $\tilde{\alpha}$, and action α is strictly better than action $\tilde{\alpha}$ for all utilities in M_{α^*}. For instance, suppose that for all utilities $v \in M_{\alpha^*}$ and for some pure actions $a \in S_{\alpha^*}$ and $\tilde{a} \in A$, we have $\tilde{a} \succ_v a$. By Proposition 2, we can conclude that α^* is not optimal.

The argument for the proof of Proposition 2 goes as follows. Take an action $\tilde{\alpha}$ with $S_{\tilde{\alpha}} \subseteq S_{\alpha^*}$ and suppose that there exists another action α that is strictly better for all utilities in M_{α^*}. Given that the set of utilities $\mathcal{W}$ is finite, the payoffs of action α^* for utilities that do not belong to M_{α^*} must be larger than the payoff of the worst-off utilities in M_{α^*} by some finite amount, say $\epsilon > 0$. Because all utilities are continuous, it is possible to add a little bit of α and subtract a little bit of $\tilde{\alpha}$ from action α^* to make all utilities in M_{α^*} better-off without rendering anyone outside M_{α^*} worse-off by more than ϵ. Notice that the resulting action is well-defined because $S_{\tilde{\alpha}} \subseteq S_{\alpha^*}$. Therefore, action α^* is not optimal.

3.1 Representation in the Marschak–Machina triangle

We conclude this section with a graphical representation of the results in Propositions 1 and 2. Figure 1 shows an example with three pure actions (a, b and c) and three utility functions (v_1, v_2 and v_3) using a revisitation of the Marschak–Machina triangle.[14] Every point in the triangle corresponds to the lottery associated with an action. The figure also includes the indifference curves for the three utilities. Given that the level of the indifference curves matters, we make it explicit through the thickness of the curves. The indifference curves of utilities v_1, v_2, and v_3 passing through $\hat{\alpha}$ have the same thickness and thus achieve the same level of utility. At the same time, the indifference curve of utility v_3 passing through α^* is thicker than the indifference curve passing through $\hat{\alpha}$ because it is associated with a higher level of utility.

[14]The canonical Marschak–Machina triangle represents the set of all lotteries involving three fixed outcomes. Here instead, we represent the set of all lotteries arising from random choices that involve three actions.

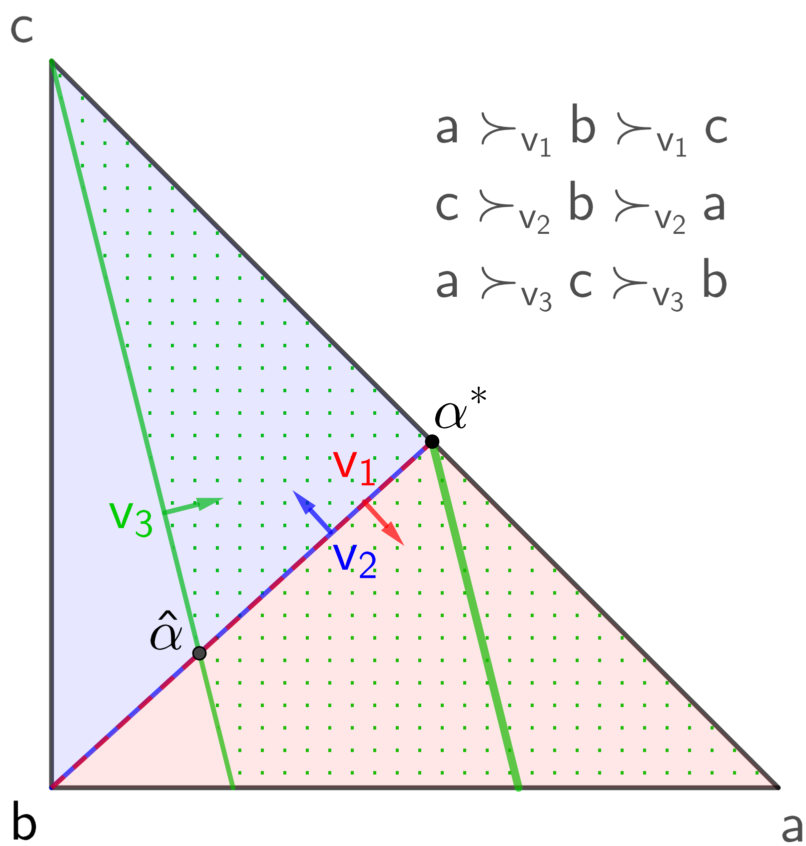

Figure 1: Example with $\mathcal{A} = \triangle(\{a, b, c\})$ and $\mathcal{W} = \{v_1, v_2, v_3\}$.

According to Proposition 1, an action α is optimal if there is no other action that is strictly better for all utilities in M_α. For instance, let us consider the mixed action $\hat{\alpha}$ and notice that $M_{\hat{\alpha}} = \{v_1, v_2, v_3\}$. Understanding whether action $\hat{\alpha}$ is optimal amounts to check whether the intersection of the strict upper contour sets of the three utilities at $\hat{\alpha}$ is empty. In Figure 1, the shaded red region is the strict upper contour set for v_1, the shaded blue region for v_2, and the dotted green region for v_3. As it is clear from the figure, the intersection is empty: to make utility v_2 strictly better-off, it is necessary to make utility v_1 strictly worse-off and vice versa. Therefore, by Proposition 1 action $\hat{\alpha}$ is optimal.

To appreciate the practical use of Proposition 2, consider a richer decision framework with the same set of utility functions but with a larger set of pure actions A'

such that $\{a, b, c\} \subset A'$. Suppose that we are interested in verifying whether an action α with support $S_\alpha = A'$ is optimal. Thanks to Proposition 2, it is still possible to address this task by looking at the Marschak–Machina triangle in Figure 1. Indeed, if there are two actions in the triangle such that for all utilities in M_α, one action is strictly better than the other, then by Proposition 2 α is not optimal.

4 Uniqueness and efficiency

Proposition 1 characterizes the set of optimal actions in terms of the strict upper contour sets of the worst-off utilities. This section explores additional properties that optimal actions might satisfy. Figure 1 shows that the set of optimal actions does not need to be a singleton. In the example, this set consists of all actions in the segment with extremes $\hat{\alpha}$ and α^*. The next proposition answers the following question: under what condition is an optimal action unique?

Proposition 3. *Let $\alpha^* \in \mathcal{A}$ be an optimal action. Then α^* is unique if and only if there is no action $\alpha \in \mathcal{A}$, with $\alpha^* \neq \alpha$, such that $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$.*

Proposition 3 states that an optimal action α^* is unique whenever there is no action α that is indifferent to α^* for all utilities in M_{α^*}. If such action exists, it is possible to mix α^* with a little bit of α. The resulting new action keeps the set of worst-off utilities fixed to M_{α^*} maintaining the same level of minimum utility. Conversely, suppose that there are two actions α^* and α that are optimal. All utilities in M_{α^*} weakly prefer action α to action α^*. Let us build a new action $\hat{\alpha}$ by mixing action α^* with a little bit of α. The resulting action $\hat{\alpha}$ is still optimal. Furthermore, the set of utilities $M_{\hat{\alpha}}$ coincides with all the utilities in M_{α^*} for which α^* is indifferent to α. Therefore, the action $\hat{\alpha}$ is optimal and all the utilities in $M_{\hat{\alpha}}$ are indifferent between actions $\hat{\alpha}$ and α^*.

If the condition in Proposition 3 fails, the set of optimal actions is not a singleton. To reduce the extent of this multiplicity, we propose an efficiency criterion that refines the set of optimal actions. The set of worst-off utilities M_{α^*} plays a key role in determining whether action α^* is optimal. For this to be the case, there must be no other action α that is strictly better than α^* for all utilities in M_{α^*}. A first natural refinement is then to ask that there is no other action α that Pareto dominates α^* in M_{α^*}. That is, there is no other action α such that α is weakly better than α^* for all utilities $v \in M_{\alpha^*}$, and α is strictly better than α^* for at least one utility $v \in M_{\alpha^*}$.

This Pareto efficiency requirement in the set of worst-off utilities relates to the following question: how large is the set of worst-off utilities? Suppose an optimal action α^* is not Pareto efficient in M_{α^*}. In this case, it is possible to find another action whose set of worst-off utilities is strictly smaller in the sense of set inclusion. The following proposition formalizes this intuition.

Proposition 4. *An optimal action α^* is Pareto efficient in M_{α^*} if and only if $M_{\alpha^*} \subseteq M_\alpha$ for any other optimal action α.*

According to Proposition 4, an optimal action α^* is Pareto efficient in M_{α^*} whenever there is no other optimal action that induces a strictly smaller set of worst-off utilities. For instance, let us come back to the scenario in Figure 1. The action $\hat{\alpha}$ is optimal because no action is strictly better for all the utilities. However, from Proposition 4, it is possible to conclude that action $\hat{\alpha}$ is not Pareto efficient in $M_{\hat{\alpha}}$. Indeed, any action α in the interval $(\hat{\alpha}, \alpha^*]$ is also optimal and $M_\alpha = \{v_1, v_2\} \subset \{v_1, v_2, v_3\} = M_{\hat{\alpha}}$. We refer to actions that are Pareto efficient in the induced set of worst-off utilities as minimal and denote by M_{min} the set of worst-off utilities that they induce.

Despite any action in the interval $(\hat{\alpha}, \alpha^*]$ is minimal, the most natural action to pick seems α^*, because utilities v_1 and v_2 are always indifferent, while utility v_3 strictly prefers action α^*. In other words, a sensible selection criterion should also impose an efficiency requirement for utilities that are outside M_{min}. This consideration leads us to our definition of efficiency.

Definition 1. *An action $\alpha^* \in \mathcal{A}$ is efficient if it is minimal and there is no other action that Pareto dominates α^* in $\mathcal{W}$.*

In the example of Figure 1, α^* is the only efficient action. The next proposition shows that there is always at least one efficient action.

Proposition 5. *For any finite set of utilities $\mathcal{W}$, there always exists an efficient action.*

The existence of a minimal action follows from Proposition 4 and by the fact that the set of utilities $\mathcal{W}$ is finite. If an optimal action is not minimal, then by Proposition 4 there must exist another optimal action that induces a strictly smaller set of worst-off utilities. Given that $\mathcal{W}$ is finite, there must exist a minimal action that induces the smallest set of worst-off utilities.

At this point, it is not possible to directly establish the existence of an efficient action by solving a maximization problem over the set of minimal actions because this set may not be compact as in the example of Figure 1. We circumvent this issue as follows. First, we maximize again ($\star$) over the set of optimal actions using $\mathcal{W} \setminus M_{min}$ as set of utility functions. Second, we show that all the actions that solve the maximization problem must be minimal. Third, within this compact subset of minimal actions, we maximize the sum of the expected utilities over all utilities in $\mathcal{W} \setminus M_{min}$. Finally, we prove that any solution to this latter maximization problem is efficient.

5 Deliberate randomization with two utilities

In this section, we study the role of deliberate randomization for a DM with convex preferences and two utilities: $\mathcal{W} = \{v_1, v_2\}$. From an operational point of view, we show that this assumption is appealing because it simplifies the structure of the set of optimal actions. At the same time, it still allows interesting deviations from expected utility. For instance, in the C-EU model, Cerreia-Vioglio et al. (2015)

show that two utilities are enough to rationalize the certainty effect in the Allais' common ratio example.

Our first result is that a DM with convex preferences and two utilities never finds it strictly beneficial to select more than two pure actions with positive probability.

Proposition 6. *If* $|\mathcal{W}| = 2$*, then*

$$\max_{\alpha \in \mathcal{A}} u(\alpha) = \max_{\alpha \in \{\alpha' \in \mathcal{A} : |S_{\alpha'}| \leq 2\}} u(\alpha).$$

There are two possibilities for any three pure actions in the support of an optimal mixed action: either both utilities are indifferent among them, or they have opposite preferences. Otherwise, the mixed action would not be optimal. In the case of indifference, it is easy to construct another optimal mixed action with smaller support. In the proof, we show that this is also possible in the scenario of opposite preferences.

To fix ideas, consider the example in Figure 1 neglecting the role of utility v_3. Utilities v_1 and v_2 have opposite preferences for the pure actions a, b and c. In particular, $a \succ_{v_1} b \succ_{v_1} c$ and $c \succ_{v_2} b \succ_{v_2} a$. We show that if there is a mixed action inside the triangle that is optimal (for instance, action $\hat{\alpha}$), then there must exist a unique mixed action α^* with support $\{a, c\}$ that is indifferent to the pure action b for both utilities. Therefore, starting from $\hat{\alpha}$, one can reduce to zero the probability weight of action b and increase the probability weights of actions a and c by $\alpha^*(a)\hat{\alpha}(b)$ and $\alpha^*(c)\hat{\alpha}(b)$, respectively. The resulting new mixed action has smaller support and is still optimal.

Proposition 6 provides a testable implication of our restriction on the set of utility functions. Experiments that document deliberate randomization typically focus on binary comparisons. For instance, in Agranov and Ortoleva (2020) subjects can use an external randomization device to choose between two lotteries, exactly as in our theoretical framework. To test our restriction on the number of utilities, one can enlarge the set of available lotteries and add a small cost for selecting more than two lotteries with a positive probability. Subjects consistent with the assumption of two utilities should never be willing to pay any positive amount.

Maintaining the assumption of two utilities, we now characterize the mixed actions that maximize the DM's preferences when there are no optimal pure actions. In this case, we call randomization strictly beneficial.

Definition 2. *Randomization is strictly beneficial if*

$$\exists \alpha \in \mathcal{A} : u(\alpha) > \max_{a \in A} u(a).$$

In what follows, we first look at the case where the DM is indifferent among all the pure actions. Then, we conclude by studying what happens when there are only two pure actions.

5.1 Indifference

A non-expected utility DM with convex preferences always strictly benefits from randomization when indifferent among all the pure actions.

Proposition 7. *Assume that* $\arg\max_{a \in A} u(a) = A$*. For any finite set of utilities* $\mathcal{W}$*, randomization is strictly beneficial if and only if there is no utility* $v \in \mathcal{W}$ *such that* $P_v = A$*.*

Whenever a utility v always belongs to the set of worst-off utilities, then it is as if the DM had expected utility preferences with utility v. This result holds regardless of the size of $\mathcal{W}$. The next proposition characterizes the set of optimal mixed actions under indifference.

Proposition 8. *Suppose that* $\mathcal{W} = \{v_1, v_2\}$*,* $\arg\max_{a \in A} u(a) = A$ *and there is no utility* $v \in \mathcal{W}$ *such that* $P_v = A$*. A mixed action* $\alpha \in \mathcal{A}$ *is optimal if and only if the following conditions hold:*

1. $M_\alpha = \{v_1, v_2\}$.

2. $S_\alpha \subseteq \underset{a \in P_{v_1} \setminus P_{v_2}}{\arg\max}\, U\left[\mathbb{E}(a, v_2), v_2\right] \cup \underset{a \in P_{v_2} \setminus P_{v_1}}{\arg\max}\, U\left[\mathbb{E}(a, v_1), v_1\right]$.

The evaluation of the optimal mixed action α must be the same for both utilities. Otherwise, it is always possible to increase the minimum evaluation. Moreover, the optimal mixed action must select with positive probability only pure actions for which the two utilities disagree in the evaluations. That is, each pure action must belong to $P_{v_1} \setminus P_{v_2}$ or $P_{v_2} \setminus P_{v_1}$. In light of Proposition 6, it is enough to consider only mixed actions that assign positive probability to two pure actions, one from each set.

Intuitively, when the two utilities have two different evaluations for a pure action, selecting the action with positive probability is strictly beneficial because it helps the DM hedging against his conservative nature. However, when the two evaluations coincide, no hedging is possible. Because of the indifference assumption, utility v_1 assigns the same value to all the actions in $P_{v_1} \setminus P_{v_2}$. Therefore, among these actions, an optimal mixed action must select only those that maximize the evaluation for utility v_2. An analogous argument applies to actions in $P_{v_2} \setminus P_{v_1}$.

5.2 Two actions

In most experiments that document deliberate randomization, there are only two feasible pure actions for each choice. The setting with binary actions is also interesting in several applications, such as the static game we consider in the next section. We begin characterizing strict benefits from randomization when there are only two pure actions.

Proposition 9. *Assume that $A = \{a, b\}$ and $\mathcal{W} = \{v_1, v_2\}$. Randomization is strictly beneficial if and only if the following are true:*

1. *There is no utility $v \in \mathcal{W}$ such that $P_v = A$.*

2. *Either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$.*

The DM can find randomization strictly beneficial even if the two pure actions do not ensure the same minimum evaluation. Therefore, the DM must be able to commit credibly to stick with the indications of the randomization device. As for the case of indifference, the DM's preferences must not degenerate to expected utility. Moreover, the two utilities must disagree in ranking the two pure actions. Randomization is a valuable tool only when this internal disagreement is present because it allows the DM to hedge against his pessimistic nature in evaluating pure actions.

Besides shedding light on the drivers that make randomization desirable, Proposition 9 also allows studying the DM's risk attitude. In the C-EU model, Cerreia-Vioglio et al. (2015) show that the DM is risk averse (respectively, risk seeking) if and only if all the utilities in $\mathcal{W}$ are concave (respectively, convex). Thanks to Proposition 9, it is possible to rule out both risk attitudes considering two pure actions a and b, where a is a mean-preserving spread of b.[15]

Corollary 1. *Assume that $\mathcal{W} = \{v_1, v_2\}$, $A = \{a, b\}$ and that action a is a mean-preserving spread of action b. If incentives to randomize are strict, then a C-EU DM is neither risk averse, nor risk seeking.*

According to Proposition 9, if incentives to randomize are strict, the two utilities must disagree in the ranking between actions a and b. But this necessarily implies that one utility is convex while the other one is concave. Consequently, a C-EU DM is neither risk averse nor risk seeking.

We conclude with the characterization of the optimal and unique mixed action.

Corollary 2. *Assume that $A = \{a, b\}$ and $\mathcal{W} = \{v_1, v_2\}$ and that randomization is strictly beneficial. The action $\alpha \in \mathcal{A}$ is uniquely optimal if and only if $M_\alpha = \mathcal{W}$.*

When there are two utilities, two actions, and incentives to randomize are strict, the unique optimal mixed action equalizes the payoff of the two utilities.

[15] We identify actions with the lotteries that they induce. Formally, given a belief μ, we say that action a is a mean-preserving spread of action b if $\hat{\rho}(a, \mu)$ is a mean-preserving spread of $\hat{\rho}(b, \mu)$. For a definition of mean-preserving spread, we refer to Rothschild and Stiglitz (1970).

6 Games with Convex Preferences

A normal game form is a mathematical structure $\langle N, (S_i)_{i\in N}, C, g\rangle$, where N is a finite set of players, S_i and is the set of available actions for each player i, C is the set of consequences and $g\colon \times_{i\in N} S_i \to C$ is the outcome function that associates consequences with strategy profiles. Therefore, each player i faces the decision framework $\langle S_i, S_{-i}, C, g\rangle$.[16]

Given a conjecture $\mu_i \in \triangle(S_{-i})$, player i chooses $\alpha_i \in \triangle(S_i)$ to maximize

$$u_i(\alpha, \mu_i) = \min_{v\in\mathcal{W}_i} \mathrm{U}\left[\mathbb{E}_{\mu_i}(\alpha, v), v\right],$$

where we make the dependence from the conjecture explicit. Similarly, we write M_{α_i,μ_i}, P_{v,μ_i}, $\succ_{v,\mu_i}$ and $\sim_{v,\mu_i}$ instead of M_{α_i}, P_v, $\succ_v$ and $\sim_v$. A normal-form game with convex preferences G adds to the normal game form the profile $(\mathcal{W}_i)_{i\in N}$ of sets of utility functions on C. Every normal-form game with convex preferences always has a Nash equilibrium because all the standard assumptions for existence hold.[17] In what follows, we characterize the set of all possible Nash equilibria when there are two players, each having two pure actions and two utility functions.

Consider a normal-form game with convex preferences G with $N = \{A, B\}$, $S_A = \{a_1, a_2\}$, $S_B = \{b_1, b_2\}$ and $|\mathcal{W}_A| = |\mathcal{W}_B| = 2$. For convenience, we identify mixed actions for players A and B with the probabilities $\alpha \in (0,1)$ and $\beta \in (0,1)$ that they assign to actions a_1 and b_1. We also denote by v and w generic utilities for players A and B. Given a utility $v \in \mathcal{W}_A$ for player A, we represent by β_v the mixed action of player B such that $a_1 \sim_{v,\beta_v} a_2$. Therefore, player A is indifferent between pure actions a_1 and a_2 when using utility v and thinking that player B chooses the mixed action β_v. Similarly, given a utility $w \in \mathcal{W}_B$ for player B, we denote by α_w the mixed action of player A such that $b_1 \sim_{w,\alpha_w} b_2$. We assume

[16]We denote by $-i = N \setminus \{i\}$ the set of players different from i.

[17]The set $\times_{i\in N} \triangle(S_i)$ is compact and convex. Moreover, for each player $i \in N$, the function u_i is continuous and quasi-concave since the preferences that it represents satisfy convexity.

that for all $v \in \mathcal{W}_A$ and $w \in \mathcal{W}_B$, $\alpha_w \in [0,1]$ and $\beta_v \in [0,1]$ are well-defined. This amounts to imposing that under no single utility, one player has a dominant action.

6.1 Strict Mixed Nash Equilibria

The profile of strategies (α_w, β_v) is the mixed Nash equilibrium that would result in a game in which players A and B maximize subjective expected utility with utilities v and w. Within the subjective expected utility framework, every player is indifferent between the mixed action played in equilibrium and all the pure actions in the support. When instead players have convex preferences, our analysis thus far shows that given a fixed conjecture about the other player's action, the incentives to play a mixed action may be strict. We now study under what conditions incentives to randomize extend to mixed Nash equilibria of G.

Let $(\alpha, \beta) \in (0,1)^2$ be a mixed Nash equilibrium of G. First, notice that

$$\alpha \in \bar{A} := [\min_{w \in \mathcal{W}_B} \alpha_w, \max_{w \in \mathcal{W}_B} \alpha_w] \text{ and } \beta \in \bar{B} := [\min_{v \in \mathcal{W}_A} \beta_v, \max_{v \in \mathcal{W}_A} \beta_v].$$

These conditions follow directly from the definition of Nash equilibrium. For instance, if $\beta > \max \bar{B}$, then either $a_1 \succ_{v,\beta} a_2$ or $a_2 \succ_{v,\beta} a_1$ for all $v \in \mathcal{W}_A$. In both cases, the mixed action $\alpha \in (0,1)$ for player A can not be a best reply to action β for player B. The next corollary clarifies that, in equilibrium, incentives to randomize are strict for both players A and B depending on whether actions β and α are boundary or interior points of $\bar{B}$ and $\bar{A}$, respectively.[18]

Corollary 3. *Let $(\alpha, \beta) \in (0,1)^2$ be a mixed Nash equilibrium of G. The following statements are true:*

1. $\alpha \in \bar{A}^{\text{o}}$ *if and only if* $u_B(\beta, \alpha) > \max\{u_B(b_1, \alpha), u_B(b_2, \alpha)\}$.
2. $\beta \in \bar{B}^{\text{o}}$ *if and only if* $u_A(\alpha, \beta) > \max\{u_A(a_1, \beta), u_A(a_2, \beta)\}$.

[18] Given a generic set X, we denote by X^{o} the set of interior points of X.

Suppose first that $\alpha \notin A^{\mathrm{o}}$. Then player B does not strictly benefit from randomization. Indeed, for any $\alpha \notin A^{\mathrm{o}}$, at least one of the two utilities in $\mathcal{W}_B$ is indifferent between actions b_1 and b_2. By Proposition 9, incentives to randomize can not be strict. Suppose instead that $\alpha \in A^{\mathrm{o}}$. In this case, player B strictly benefits from randomization. Given that $\alpha \in A^{\mathrm{o}}$, it must be that one utility in $\mathcal{W}_B$ strictly prefers action b_1 to action b_2, and the other has opposite preferences. By Proposition 9, it is enough to show that there is no utility in $\mathcal{W}_B$ that belongs to the sets of worst-off utilities induced by both pure actions. This latter condition must hold because otherwise, the mixed action β would not be a best reply to the correct conjecture α. An analogous reasoning can be used to prove the second statement of Corollary 3.

In light of Corollary 3, we classify mixed Nash equilibria as follows.

Definition 3. *Let $(\alpha, \beta) \in (0,1)^2$ be a mixed Nash equilibrium of G. We call (α, β)*

- *weak if $\alpha \notin \bar{A}^{\mathrm{o}}$ and $\beta \notin \bar{B}^{\mathrm{o}}$.*
- *partially strict if either $\alpha \notin \bar{A}^{\mathrm{o}}$ and $\beta \in \bar{A}^{\mathrm{o}}$ or $\alpha \in \bar{A}^{\mathrm{o}}$ and $\beta \notin \bar{A}^{\mathrm{o}}$.*
- *strict if $\alpha \in \bar{A}^{\mathrm{o}}$ and $\beta \in \bar{A}^{\mathrm{o}}$.*

When the sets $\bar{A}^{\mathrm{o}}$ and $\bar{B}^{\mathrm{o}}$ are empty, Corollary 3 implies that there can not be strict mixed Nash equilibria. One notable example of a game in which this happens is matching pennies. Indeed, as long as the utilities of the two players are strictly increasing, we have $\bar{A} = \bar{B} = \{0.5\}$. For the remaining part of this section, we assume that $\bar{A}^{\mathrm{o}}$ and $\bar{B}^{\mathrm{o}}$ are non-empty so that any mixed Nash equilibrium is a priori possible.

The computation of weak equilibria follows the same logic used to compute equilibria under expected utility. In equilibrium, each player must be indifferent between the two pure actions. We now turn to the analysis of partially strict and strict equilibria, in which at least one player strictly benefits from randomization. In particular, let us focus on player A and suppose that there exists a subset of conjectures $X \subseteq \bar{B}^{\mathrm{o}}$ under which player A strictly benefits from randomization.

			B	
			b_1	b_2
A	a_1		3,1	0,0
	a_2		0,0	1,3

Table 1: Coordination game: outcome function.

The next corollary characterizes the best reply of player A for all conjectures in X, assuming the maxmin expected utility model.

Corollary 4. *Let $W_A = \{v_A, w_A\}$. For all conjectures $\beta \in X$, the unique optimal mixed action $\alpha(\beta)$ in the maxmin expected utility model satisfies*

$$\alpha(\beta) = \frac{\mathbb{E}_\beta[a_2, w_A] - \mathbb{E}_\beta[a_2, v_A]}{\mathbb{E}_\beta[a_2, w_A] - \mathbb{E}_\beta[a_2, v_A] + \mathbb{E}_\beta[a_1, v_A] - \mathbb{E}_\beta[a_1, w_A]}.$$

Corollary 4 provides a simple closed-form expression that characterizes the best reply of player A for the subset of conjectures under which randomization is strictly beneficial. An interesting insight that emerges from Corollary 4 is that the optimal probability with which player A chooses action a_1 is increasing in $|\mathbb{E}_\beta[a_2, w_A] - \mathbb{E}_\beta[a_2, v_A]|$ and decreasing in $|\mathbb{E}_\beta[a_1, w_A] - \mathbb{E}_\beta[a_1, v_A]|$. That is, in the maxmin expected utility model, players dislike actions for which there is high variability in their evaluations.

To illustrate all the possible types of mixed Nash equilibria, we consider the coordination game with the outcome function represented in Table 1. We assume that players A and B behave according to maxmin expected utility criterion. Each player has two utility functions, one CARA and one CRRA.[19] Figure 2 represents the best replies for the two players.[20] Every intersection point of the two best replies represents a Nash equilibrium. In this example, there are 11 Nash equilibria: two pure, four weak, four partially strict, and one strict. The two pure

[19]CARA: $v_A(x) = w_A(x) = 1 - \frac{1}{\alpha}e^{-\alpha x}$, with $x \geq 0$ and $\alpha > 0$. CRRA: $v_B(x) = w_B(x) = x^\gamma$, with $x \geq 0$ and $\gamma \in (0, 1)$.

[20]Parameters: $\alpha = 1.52$ and $\gamma = 0.42$.

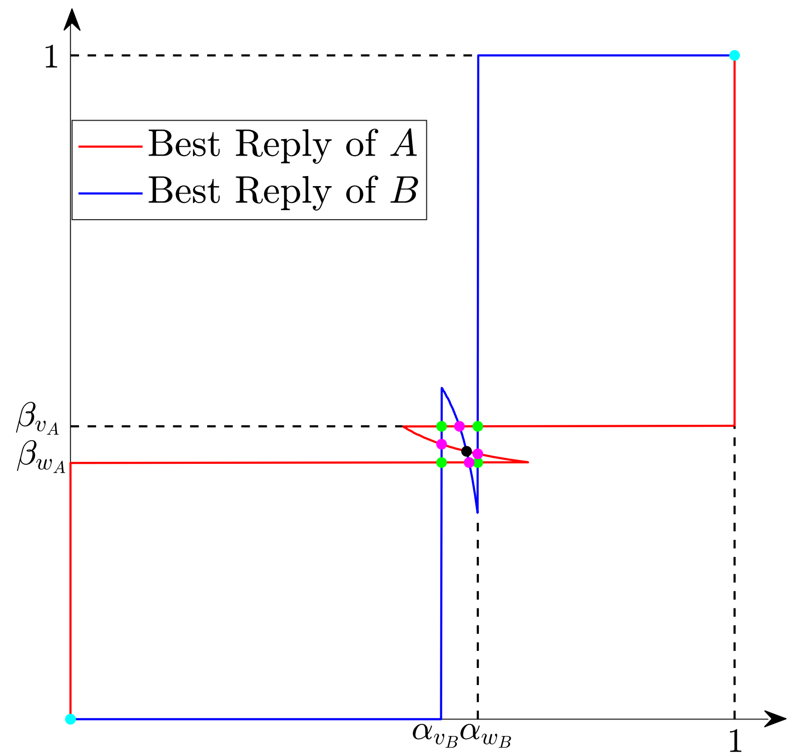

Figure 2: Nash equilibria of the coordination game.

Nash equilibria are in light blue, the four weak Nash equilibria are in light green, the four partially strict mixed Nash equilibria are in magenta, and the strict Nash equilibrium is in black.

In this simple and analytically tractable scenario of two utility functions for each player, we obtain starkly different predictions from the expected utility case. From a numerical point of view, convexity may lead to a multiplicity of mixed equilibria. Most importantly, partially strict and strict mixed Nash equilibria do not have an analog under expected utility. We now extend the notion of efficiency developed in Section 4 to profiles of actions and show that strict mixed Nash equilibria are the only type of equilibria that satisfy this notion.

Definition 4. *A mixed Nash equilibrium* (α, β) *of* G *is efficient if in equilibrium* α *and* β *are efficient.*

The next corollary clarifies that only strict mixed Nash equilibria satisfy the efficiency requirement in Definition 4.

Corollary 5. *Let* (α, β) *be a mixed Nash equilibrium of G and suppose that* $\bar{A}^{o}$ *and* $\bar{B}^{o}$ *are non-empty. Then* (α, β) *is efficient if and only if* (α, β) *is strict.*

The notion of efficiency introduced in Definition 4 can serve as a selection criterion for settings with multiple mixed Nash equilibria. In the example described in Figure 2, there are nine mixed Nash equilibria, but only one is strict and, by Corollary 5, efficient.

Overall, the presence of strict mixed Nash equilibria is the main element of novelty that emerges from our equilibrium analysis in games with convex preferences. We show that, as documented by contemporaneous experimental works, incentives to randomize extend to strategic interaction settings but only under certain conditions. For instance, strict mixed Nash equilibria may arise in a coordination game as the one described in Figure 2 but do not exist in matching pennies. Besides being empirically relevant given the evidence of randomization in games, Corollary 5 shows that when they exist, strict mixed Nash equilibria are also normatively appealing because they are the only efficient equilibria.

7 Conclusions

Despite the growing theoretical and experimental literature on random choices under risk, the applicability of models that rationalize deliberate randomization is still limited. This paper studies the set optimal actions for a DM whose preferences satisfy convexity, the axiom that makes randomization weakly beneficial. Under convexity, the DM's preferences admit a conservative multi-utility representation: actions are ranked only through the lowest utility valuation they generate.

One drawback of this representation in applications is that it is not differentiable, so standard optimization techniques are not viable. Our main result (Proposition 1) shows that an action is optimal whenever the intersection of the strict upper contour sets of the worst-off utilities is empty. When more than one action is optimal, we propose Pareto efficiency in the set of worst-off utilities and in the set of all utilities as a selection rule. Proposition 4 clarifies that the first requirement amounts to isolating optimal actions that induce the smallest set of worst-off utilities. The second requirement allows instead to account for utilities outside this set. Proposition 5 guarantees that there is always an optimal action that satisfies both requirements and thus is efficient.

Next, we narrow our attention to random choices for a DM with two utilities. Proposition 6 provides a testable implication of this assumption, proving that a DM with two utilities never finds it strictly beneficial to select more than two actions with positive probability. We then study under what conditions randomization is strictly beneficial and the properties that an optimal random choice must satisfy in two cases: when the DM is indifferent among the pure actions (Proposition 7 and Proposition 8) and when there are only two pure actions (Proposition 9 and Corollary 2).

The binary actions setting recreates the typical environment that subjects face in experiments on randomization under risk. In general, preferences for randomization can coexist with various attitudes towards risk. Our analysis of randomization incentives suggests a new approach to rule out risk aversion and risk seeking in the C-EU model. According to Corollary 1, a C-EU DM is neither risk averse nor risk seeking if incentives to randomize are strict between two actions, one being a mean preserving spread of the other. Moreover, Corollary 2 shows that when incentives to randomize are strict, the optimal mixed action is unique.

A special case of the decision framework that we study is game theory. We focus on a generic game with two players, each with two actions and two utility functions. The new prediction that arises from our analysis is that strict incentives to randomize extend to strategic interaction settings. Corollary 3 provides necessary conditions for the existence of a new class of mixed Nash equilibria that we call strict because players strictly prefer the equilibrium mixed actions to the pure actions. Corollary 4 derives a closed-form expression of the best response function for the case in which randomization is strictly beneficial, and players have maxmin preferences. We then exploit this result to compute the mixed Nash equilibria of a simple coordination game. In this example, we find nine mixed Nash equilibria, one of which is strict. Although convexity may lead to a multiplicity of mixed equilibria, we show in Corollary 5 that when they exist, only strict equilibria are such that all the mixed actions are efficient.

Appendix: proofs

This appendix contains the proofs of the results stated in the body of the text.

Proof of Proposition 1

Step 1. *If* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha^*) = \emptyset$*, then* α^* *is optimal.*

Proof of Step 1. Consider an action α^* such that $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha^*) = \emptyset$. This implies that for all $\alpha \in \mathcal{A}$, there exists $v \in M_{\alpha^*}$ such that

$$u(\alpha^*) = \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right] \geq \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha).$$

Therefore, α^* is optimal.

Step 2. *If* α^* *is optimal, then* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha^*) = \emptyset$*.*

Proof of Step 2. We show that if there exists an action $\alpha \in \mathcal{A}$ such that $\alpha \succ_v \alpha^*$ for all $v \in M_{\alpha^*}$, then α^* is not optimal. Define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0,1)$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \lambda\alpha(a) + (1-\lambda)\alpha^*(a).$$

We now show that there exists a value $\bar{\lambda} \in (0,1]$ such that for all $\lambda \in (0, \bar{\lambda})$, $u(\hat{\alpha}_\lambda) > u(\alpha^*)$. To this end, consider the function $\Psi\colon [0,1] \times \mathcal{W} \to \mathbb{R}$ such that $\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*)$, for all $\lambda \in [0,1]$ and $v \in \mathcal{W}$. As an intermediate step, we prove that for each $v \in \mathcal{W}$, there exists a value $\bar{\lambda}_v \in (0,1]$ such that for all $\lambda \in (0, \bar{\lambda}_v)$, $\Psi(\lambda, v) > 0$. Take a utility $v \in \mathcal{W}$. There are two possibilities:

1. If $\alpha \succsim_v \alpha^*$, then for all $\lambda \in (0,1]$

$$\begin{aligned}
\Psi(\lambda, v) &= \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) \\
&= \mathrm{U}\left[\lambda\mathbb{E}(\alpha, v) + (1-\lambda)\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \\
&\geq \mathrm{U}\left[\lambda\mathbb{E}(\alpha^*, v) + (1-\lambda)\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \\
&= \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \geq 0,
\end{aligned}$$

where at least one of the two weak inequalities holds strict. If $\alpha \succ_v \alpha^*$, the first inequality is strict because the function $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. If instead $\alpha \sim_v \alpha^*$, the last inequality is strict because $v \notin M_{\alpha^*}$. For all such v, we let $\bar{\lambda}_v = 1$.

2. If $\alpha^* \succ_v \alpha$, then $\Psi(\cdot, v)$ is strictly decreasing in the first argument because $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. $\Psi(\cdot, v)$ is continuous in the first argument because $\mathrm{U}[\cdot, v]$ is continuous in the first argument. If $\Psi(1, v) \geq 0$, the result follows immediately by taking $\bar{\lambda}_v = 1$. Suppose that $\Psi(1, v) < 0$. Notice that $\Psi(0, v) > 0$ because $\alpha^* \succ_v \alpha$ implies that $v \notin M_{\alpha^*}$. Therefore, by the Intermediate Value Theorem, there exists $\bar{\lambda}_v \in (0, 1)$ such that $\Psi(\bar{\lambda}_v, v) = 0$. By $\Psi(\cdot, v)$ strictly decreasing, $\Psi(\lambda, v) > 0$ for all $\lambda \in (0, \bar{\lambda}_v)$.

To conclude the proof, we let $\bar{\lambda} = \min_{v \in \mathcal{W}} \bar{\lambda}_v$. Given that $\mathcal{W}$ is finite, $\bar{\lambda}$ is well-defined. By construction, for all $\lambda \in (0, \bar{\lambda})$ and for all $v \in \mathcal{W}$,

$$\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) > 0.$$

This implies that $u(\hat{\alpha}_\lambda) > u(\alpha^*)$. Consequently, we conclude that α^* is not optimal.

Proof of Proposition 2

Step 1. *If* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha) = \emptyset$ *for all* $\alpha \in \mathcal{A}$ *with* $S_\alpha \subseteq S_{\alpha^*}$*, then* α^* *is optimal.*

Proof of Step 1. Consider an action $\alpha^* \in \mathcal{A}$. If for all $\alpha \in \mathcal{A}$ with $S_\alpha \subseteq S_{\alpha^*}$, $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha) = \emptyset$, then trivially α^* is optimal by Proposition 1.

Step 2. *If* α^* *is optimal, then* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha) = \emptyset$ *for all* $\alpha \in \mathcal{A}$ *with* $S_\alpha \subseteq S_{\alpha^*}$*.*

Proof of Step 2. Take an action $\alpha^* \in \mathcal{A}$. We show that if there exist an action $\tilde{\alpha} \in \mathcal{A}$, with $S_{\tilde{\alpha}} \subseteq S_{\alpha^*}$ and another action $\alpha \in \mathcal{A}$ such that $\alpha \succ_v \tilde{\alpha}$ for all $v \in M_{\alpha^*}$, then α^* is not optimal. Define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0, \hat{\lambda})$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \alpha^*(a) + \lambda\left[\alpha(a) - \tilde{\alpha}(a)\right],$$

where the upper bound $\hat{\lambda}$ is defined as follows:

$$\hat{\lambda} = \max \{\lambda \in (0,1] : \forall a \in A, \hat{\alpha}_\lambda(a) \geq 0\}.$$

Notice that $\hat{\lambda}$ is well-defined because $S_{\tilde{\alpha}} \subseteq S_{\alpha^*}$. We now show that there exists a value $\bar{\lambda} \in (0, \hat{\lambda}]$ such that for all $\lambda \in (0, \bar{\lambda})$, $u(\hat{\alpha}_\lambda) > u(\alpha^*)$. To this end, consider the function $\Psi \colon [0, \hat{\lambda}] \times \mathcal{W} \to \mathbb{R}$ such that $\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*)$, for all $\lambda \in [0, \hat{\lambda}]$ and $v \in \mathcal{W}$. As an intermediate step, we prove that for each $v \in \mathcal{W}$, there exists a value $\bar{\lambda}_v \in (0, \hat{\lambda}]$ such that for all $\lambda \in (0, \bar{\lambda}_v)$, $\Psi(\lambda, v) > 0$. Take a utility $v \in \mathcal{W}$. There are two possibilities:

1. If $\alpha \succsim_v \tilde{\alpha}$, then for all $\lambda \in (0, \hat{\lambda}]$

$$\begin{aligned} \Psi(\lambda, v) &= \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) \\ &= \mathrm{U}\left[\mathbb{E}(\alpha^*, v) + \lambda\left[\mathbb{E}(\alpha, v) - \mathbb{E}(\tilde{\alpha}, v)\right], v\right] - u(\alpha^*) \\ &\geq \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \geq 0, \end{aligned}$$

   where at least one of the two weak inequalities holds strict. If $\alpha \succ_v \tilde{\alpha}$, the first inequality is strict because the function $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. If instead $\alpha \sim_v \tilde{\alpha}$, the last inequality is strict because $v \notin M_{\alpha^*}$. For all such v, we let $\bar{\lambda}_v = \hat{\lambda}$.

2. If $\tilde{\alpha} \succ_v \alpha$, then $\Psi(\cdot, v)$ is strictly decreasing in the first argument because $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. $\Psi(\cdot, v)$ is continuous in the first argument because $\mathrm{U}[\cdot, v]$ is continuous in the first argument. If $\Psi(\hat{\lambda}, v) \geq 0$, the result follows immediately by taking $\bar{\lambda}_v = \hat{\lambda}$. Suppose that $\Psi(\hat{\lambda}, v) < 0$. Notice that $\Psi(0, v) > 0$ because $\tilde{\alpha} \succ_v \alpha$ implies that $v \notin M_{\alpha^*}$. Therefore, by the Intermediate Value Theorem, there exists $\bar{\lambda}_v \in (0, \hat{\lambda})$ such that $\Psi(\bar{\lambda}_v, v) = 0$. By $\Psi(\cdot, v)$ strictly decreasing, $\Psi(\lambda, v) > 0$ for all $\lambda \in (0, \bar{\lambda}_v)$.

To conclude the proof, we let $\bar{\lambda} = \min_{v \in \mathcal{W}} \bar{\lambda}_v$. Given that $\mathcal{W}$ is finite, $\bar{\lambda}$ is well-defined.

By construction, for all $\lambda \in (0, \bar{\lambda})$ and for all $v \in \mathcal{W}$,

$$\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) > 0.$$

This implies that $u(\hat{\alpha}_\lambda) > u(\alpha^*)$ and α^* is not optimal.

Proof of Proposition 3

Suppose that $\alpha^* \in \mathcal{A}$ is optimal.

Step 1. *If there is no $\alpha \in \mathcal{A}$ with $\alpha^* \neq \alpha$, such that $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$, then α^* is unique.*

Proof of Step 1. We show that if α^* is not unique, then there is an action $\alpha \in \mathcal{A}$, with $\alpha^* \neq \alpha$, such that $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$. Take an optimal action α with $\alpha \neq \alpha^*$. First, notice that for all $v \in M_{\alpha^*}$, it holds that $\alpha \succsim_v \alpha^*$. Consider the set $\{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\}$. By Proposition 1, this set is non-empty. If it coincides with M_{α^*}, then the proof is completed: $\alpha \sim_v \alpha^*$ for all $v \in M_{\alpha^*}$. Otherwise, define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0, 1)$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \lambda\alpha(a) + (1 - \lambda)\alpha^*(a).$$

We now show that there exists a value $\bar{\lambda} \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda})$, $\hat{\alpha}_\lambda$ is optimal and for all $v \in M_{\hat{\alpha}_\lambda}$, $\hat{\alpha}_\lambda \sim_v \alpha^*$. To this end, consider the function $\Psi \colon [0, 1] \times \mathcal{W} \to \mathbb{R}$ such that $\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*)$, for all $\lambda \in [0, 1]$ and $v \in \mathcal{W}$. As an intermediate step, we prove that for each $v \in \mathcal{W}$, there exists a value $\bar{\lambda}_v \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda}_v)$, $\Psi(\lambda, v) \geq 0$, with equality holding only if $v \in \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\}$. Take a utility $v \in \mathcal{W}$. There are two possibilities:

1. If $\alpha \succsim_v \alpha^*$, then for all $\lambda \in (0, 1]$

$$\begin{aligned}\Psi(\lambda, v) &= \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) \\ &= \mathrm{U}\left[\lambda\mathbb{E}(\alpha, v) + (1 - \lambda)\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \\ &\geq \mathrm{U}\left[\lambda\mathbb{E}(\alpha^*, v) + (1 - \lambda)\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \\ &= \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \geq 0,\end{aligned}$$

where both weak inequalities hold as equalities only if $v \in \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\}$. For all v with $\alpha \succsim_v \alpha^*$, we let $\bar{\lambda}_v = 1$.

2. If $\alpha^* \succ_v \alpha$, then $\Psi(\cdot, v)$ is strictly decreasing in the first argument because $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. $\Psi(\cdot, v)$ is continuous in the first argument because $\mathrm{U}[\cdot, v]$ is continuous in the first argument. If $\Psi(1, v) \geq 0$, the result follows immediately by taking $\bar{\lambda}_v = 1$. Suppose that $\Psi(1, v) < 0$. Notice that $\Psi(0, v) > 0$ because $\alpha^* \succ_v \alpha$ implies that $v \notin M_{\alpha^*}$. Therefore, by the Intermediate Value Theorem, there exists $\bar{\lambda}_v \in (0, 1)$ such that $\Psi(\bar{\lambda}_v, v) = 0$. By $\Psi(\cdot, v)$ strictly decreasing, $\Psi(\lambda, v) > 0$ for all $\lambda \in (0, \bar{\lambda}_v)$.

To conclude the proof, we let $\bar{\lambda} = \min_{v \in \mathcal{W}} \bar{\lambda}_v$. Given that $\mathcal{W}$ is finite, $\bar{\lambda}$ is well-defined. By construction, for all $\lambda \in (0, \bar{\lambda})$ and for all $v \in \mathcal{W}$,

$$\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) \geq 0,$$

where the inequality holds as an equality for $v \in \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\}$. Therefore, $u(\alpha^*) = u(\hat{\alpha}_\lambda)$ and given that action α^* is optimal, action $\hat{\alpha}_\lambda$ must be optimal as well. Moreover, by construction, $M_{\hat{\alpha}_\lambda} = \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\} \subset M_{\alpha^*}$ and consequently $\hat{\alpha}_\lambda \sim_v \alpha^*$ for all $v \in M_{\hat{\alpha}_\lambda}$.

Step 2. *If α^* is unique, then there is no $\alpha \in \mathcal{A}$ with $\alpha^* \neq \alpha$, such that $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$.*

Proof of Step 2. Suppose that there exist two actions α^* and α, with $\alpha^* \neq \alpha$, such that α^* is optimal and $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$. We show that α^* is not unique. To this end, define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0, 1)$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \lambda \alpha(a) + (1 - \lambda)\alpha^*(a).$$

Replicating the same argument used for Step 1 of Proposition 3, we can show that there exists a value $\bar{\lambda} \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda})$, $\hat{\alpha}_\lambda$ is optimal and $M_{\hat{\alpha}_\lambda} = M_{\alpha^*}$. This proves that the optimal action α^* is not unique.

Proof of Proposition 4

Step 1. *For any optimal action α^*, if $M_{\alpha^*} \subseteq M_\alpha$ for any other optimal action α, then α^* is efficient.*

Proof of Step 1. Take an action α^* that is optimal. We show that if α^* is not efficient, then there exists another optimal action $\hat{\alpha}$ such that $M_{\hat{\alpha}} \subset M_{\alpha^*}$. Given that α^* is not efficient, there exists $\alpha \in \mathcal{A}$ such that $\alpha \succsim_v \alpha^*$ for all $v \in M_{\alpha^*}$ and $\alpha \succ_v \alpha^*$ for some $v \in M_{\alpha^*}$. Also notice that by Proposition 1 there exists at least one $v \in M_{\alpha^*}$ with $\alpha \sim_v \alpha^*$ because α^* is optimal. Define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0,1)$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \lambda\alpha(a) + (1-\lambda)\alpha^*(a).$$

Replicating the same argument used for Step 1 of Proposition 3, we can show that there exists a value $\bar{\lambda} \in (0,1]$ such that for all $\lambda \in (0,\bar{\lambda})$, $\hat{\alpha}_\lambda$ is optimal and $M_{\hat{\alpha}_\lambda} = \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\} \subset M_{\alpha^*}$.

Step 2: *If α^* is efficient, then $M_{\alpha^*} \subseteq M_\alpha$ for any other optimal action α.*

Proof of Step 2. Take an optimal action α^*. We show that if there exists another optimal action α such that it is not true that $M_{\alpha^*} \subseteq M_\alpha$, then action α^* is not efficient. First, notice that for all $v \in M_{\alpha^*}$,

$$\mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha) = u(\alpha^*) = \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right].$$

In particular, there is a utility $\bar{v} \in M_{\alpha^*}$ such that $\bar{v} \notin M_\alpha$. For this utility $\bar{v}$, the weak inequality holds strict. Therefore, α^* is not efficient.

Proof of Proposition 5

Step 1. *There exists an optimal action α^* that is Pareto efficient in M_{α^*}.*

Proof of Step 1. By contradiction, suppose that there is no optimal action α^* that is Pareto efficient in M_{α^*}. Take an optimal action α. By Proposition 4, there exists another optimal action $\hat{\alpha}$ such that $M_{\hat{\alpha}} \subset M_\alpha$. Given that $\mathcal{W}$ is finite, after a finite number of iterations of this argument, we obtain an optimal action α^* such that $M_{\alpha^*} = \{v\}$, for some utility $v \in \mathcal{W}$. By assumption, α^* is not Pareto efficient in M_{α^*}. This implies that there exists another action $\tilde{\alpha} \succ_v \alpha^*$. By Proposition 1, α^* is not optimal, a contradiction. Consequently, the set of minimal actions is non-empty. Before proceeding with Step 2 of the proof, we introduce some additional notation.

Let us denote by O the set of optimal actions, by $E \subseteq O$ the set of minimal actions and by M_{min} the set of worst-off utilities induced by minimal actions. By Step 1, the sets E and M_{min} are well-defined. Let us consider the following optimization problem:

$$\max_{\alpha \in O} \min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right]. \tag{1}$$

The set of optimal actions O is non-empty and compact by the maximum theorem. Moreover, the objective function is continuous and quasi-concave. Consequently, the set of solutions for problem (1) that we denote by E_1 is non-empty and again compact by the maximum theorem.

Step 2: *$E_1 \subseteq E$.*

Proof of Step 2. Take any optimal action α that is not minimal. We show that $\alpha \notin E_1$. Take any other action $\tilde{\alpha} \in E$. We have

$$\begin{aligned}
\min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] &= \min_{v \in \mathcal{W}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \\
&= \min_{v \in \mathcal{W}} \mathrm{U}\left[\mathbb{E}(\tilde{\alpha}, v), v\right] \\
&< \min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\tilde{\alpha}, v), v\right],
\end{aligned}$$

where the first equality holds because α is not minimal and consequently $M_{min} \subset M_{\alpha}$, the second equality holds because both actions α and $\tilde{\alpha}$ are optimal, the third strict inequality holds because $\tilde{\alpha}$ is efficient and consequently $M_{\tilde{\alpha}} = M_{min}$.

Consider now the following maximization problem:

$$\max_{\alpha \in E_1} \sum_{v \in \mathcal{W} \setminus M_{min}} \mathbb{E}(\alpha, v). \tag{2}$$

We denote by E_2 the set of actions that solve (2). By the maximum theorem, E_2 is non-empty and compact.

Step 3: *Any $\alpha \in E_2$ is efficient.*

Proof of Step 3. Take any actions $\alpha \in E_2$ and $\tilde{\alpha} \in \mathcal{A}$. If $\tilde{\alpha} \notin O$, then given that α is optimal, there must exist a utility $v \in \mathcal{W}$ such that $\alpha \succ_v \tilde{\alpha}$. If instead $\tilde{\alpha} \in O \setminus E$, then $M_{\tilde{\alpha}} \setminus M_{min} \neq \emptyset$. For any utility $v \in M_{\tilde{\alpha}} \setminus M_{min}$, we have $\alpha \succ_v \tilde{\alpha}$. Let us assume now that $\tilde{\alpha} \in E$. Notice that $E_2 \subseteq E_1 \subseteq E$. We consider two different cases:

- **Case 1**: $\tilde{\alpha} \in E_1$. Given that $\tilde{\alpha} \in E_1$ and $\alpha \in E_2$, we have

$$\sum_{v \in \mathcal{W} \setminus M_{min}} \mathbb{E}(\alpha, v) \geq \sum_{v \in \mathcal{W} \setminus M_{min}} \mathbb{E}(\tilde{\alpha}, v).$$

  Consequently, if there exists a utility $v \in \mathcal{W} \setminus M_{min}$ such that $\tilde{\alpha} \succ_v \alpha$, then there must also exist a utility $\bar{v} \in \mathcal{W} \setminus M_{min}$ such that $\alpha \succ_{\bar{v}} \tilde{\alpha}$. Therefore, no action $\tilde{\alpha} \in E_1$ Pareto dominates $\alpha \in E_2$ in $\mathcal{W}$.

- **Case 2**: $\tilde{\alpha} \in E \setminus E_1$. In this case, for at least one utility $\bar{v} \in \mathcal{W} \setminus M_{min}$, we have

$$\begin{aligned} \mathrm{U}\left[\mathbb{E}(\alpha, \bar{v}), \bar{v}\right] &\geq \min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \\ &> \min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\tilde{\alpha}, v), v\right] \\ &= \mathrm{U}\left[\mathbb{E}(\tilde{\alpha}, \bar{v}), \bar{v}\right], \end{aligned}$$

  where the first weak inequality hols for any utility $\bar{v} \in \mathcal{W} \setminus M_{min}$, the second strict inequality holds because $\alpha \in E_1$ and $\tilde{\alpha} \notin E_1$, the third equality holds for at least one utility $\bar{v} \in \mathcal{W} \setminus M_{min}$ because $\mathcal{W}$ is finite. For such utility $\bar{v}$, we have $\alpha \succ_v \tilde{\alpha}$. Therefore, no action $\tilde{\alpha} \in E \setminus E_1$ Pareto dominates $\alpha \in E_2$ in $\mathcal{W}$.

Proof of Proposition 6

Let $\mathcal{W} = \{v_1, v_2\}$ and suppose that the action $\alpha \in \mathcal{A}$ with $|S_\alpha| > 2$ is optimal. We show that there exists another optimal action $\hat{\alpha} \in \mathcal{A}$ with $|S_{\hat{\alpha}}| \leq 2$. If $u(\alpha) = u(a)$ for some $a \in A$, the statement follows. Suppose that $u(\alpha) > u(a)$ for all $a \in A$.

Step 1. *If α is optimal, then $M_\alpha = \{v_1, v_2\}$.*

Proof of Step 1. Suppose that $|M_\alpha| = 1$. Without loss of generality, let $M_\alpha = \{v_1\}$. Given that α is optimal, by Proposition 1 $\alpha \succsim_{v_1} a$ for all $a \in S_\alpha$. Therefore, $\alpha \sim_{v_1} a$ for all $a \in S_\alpha$. Given that $u(\alpha) > u(a)$ for all $a \in A$, it must be that $M_a = \{v_2\}$ for all $a \in S_\alpha$. Therefore, we have

$$\mathrm{U}\left[\mathbb{E}(\alpha, v_1), v_1\right] = \mathrm{U}\left[\max_{\tilde{a} \in S_\alpha} \mathbb{E}(\tilde{a}, v_1), v_1\right] > \mathrm{U}\left[\max_{\tilde{a} \in S_\alpha} \mathbb{E}(\tilde{a}, v_2), v_2\right] \geq \mathrm{U}\left[\mathbb{E}(\alpha, v_2), v_2\right],$$

where the equality holds because $\alpha \sim_{v_1} a$ for all $a \in S_\alpha$, the strict inequality because $M_a = \{v_2\}$ for all $a \in S_\alpha$ and the weak inequality because $\mathrm{U}[\cdot, v_2]$ is a strictly increasing function. Consequently, $\mathrm{U}[\mathbb{E}(\alpha, v_1), v_1] > \mathrm{U}[\mathbb{E}(\alpha, v_2), v_2]$, which contradicts $M_\alpha = \{v_1\}$.

Step 2. *If α is optimal, for all $a, a' \in S_\alpha$, we have $a \sim_{v_1} a'$ if and only if $a \sim_{v_2} a'$.*

Proof of Step 2. Suppose that there exist actions $a, a' \in S_\alpha$ such that $a \sim_{v_1} a'$ and $a \succ_{v_2} a'$. We show that α is not optimal. Consider a new action $\tilde{\alpha}$ such that $\tilde{\alpha}(a) = \alpha(a) + \alpha(a')$, $\tilde{\alpha}(a') = 0$ and $\tilde{\alpha}(a'') = \alpha(a'')$ for all $a'' \in S_\alpha \setminus \{a, a'\}$. Therefore, we have $\tilde{\alpha} \sim_{v_1} \alpha$ and $\tilde{\alpha} \succ_{v_2} \alpha$. Given that by Step 1 $M_\alpha = \{v_1, v_2\}$, then $M_{\tilde{\alpha}} = \{v_1\}$ and $u(\alpha) = u(\tilde{\alpha})$. However, Step 1 also implies that $\tilde{\alpha}$ is not optimal, which in turn implies that α is not optimal.

Whenever there are two actions $a, a' \in S_\alpha$, such that $a \sim_{v_1} a'$ and $a \sim_{v_2} a'$, it is possible to construct another optimal mixed action α' with $S_{\alpha'} \subset S_\alpha$. It is enough take α' such that $\alpha'(a) = \alpha(a) + \alpha(a')$, $\alpha'(a') = 0$ and $\alpha'(a'') = \alpha(a'')$ for all $a'' \in S_\alpha \setminus \{a, a'\}$. From now on, assume that there are no actions $a, a' \in S_\alpha$, such that $a \sim_{v_1} a'$ and $a \sim_{v_2} a'$, but still $|S_\alpha| > 2$.

Step 3. *If α is optimal, for all $a, a' \in S_\alpha$, we have $a \succ_{v_1} a'$ if and only if $a' \succ_{v_2} a$.*

Proof of Step 3. Suppose that there exist actions $a, a' \in S_\alpha$ such that $a \succ_{v_1} a'$ and $a \succsim_{v_2} a'$. We show that α is not optimal. By Step 2 we have $a \succ_{v_2} a'$, otherwise we could conclude $a \sim_{v_1} a'$. Therefore, by Proposition 2 we conclude that α is not optimal.

Step 4. *If α is optimal and $|S_\alpha| = n > 2$, there exists another optimal action $\tilde{\alpha}$ with $|S_{\tilde{\alpha}}| = n - 1$.*

Proof of Step 4. Consider three actions $a, a', a'' \in S_\alpha$. Without loss of generality, assume $a \succ_{v_1} a' \succ_{v_1} a''$. By Step 3, $a'' \succ_{v_2} a' \succ_{v_2} a$. Define a new mixed action α_λ parametrized by $\lambda \in [0,1]$ such that $\alpha_\lambda(a) = \lambda$ and $\alpha_\lambda(a'') = 1 - \lambda$. Consider the sets

$$S_1 := \{\lambda \in [0,1] : \alpha_\lambda \succsim_{v_1} a'\} \quad \text{and} \quad S_2 := \{\lambda \in [0,1] : \alpha_\lambda \succsim_{v_2} a'\}.$$

These sets are non-empty because $1 \in S_1$ and $0 \in S_2$. Moreover, they are closed because $\succsim_{v_1}$ and $\succsim_{v_2}$ are expected utility preferences and satisfy continuity. Therefore, let $k_1, k_2 \in (0,1)$ such that $S_1 = [k_1, 1]$ and $S_2 = [0, k_2]$. We have that $\alpha_{k_1} \sim_{v_1} a'$ and $\alpha_{k_2} \sim_{v_2} a'$. In what follows, we show that if α is optimal, then $k_1 = k_2$. If $k_1 < k_2$, for $\lambda \in (k_1, k_2)$ we have $\alpha_\lambda \succ_{v_1} a'$ and $\alpha_\lambda \succ_{v_2} a'$. By Proposition 2, α is not optimal. If $k_1 > k_2$, for $\lambda \in (k_2, k_1)$ we have $a' \succ_{v_1} \alpha_\lambda$ and $a' \succ_{v_2} \alpha_\lambda$. By Proposition 2, α is not optimal. Therefore, it must be that $k_1 = k_2 = k$. Define a new mixed action $\hat{\alpha}$ such that

- $\hat{\alpha}(a') = 0$.
- $\hat{\alpha}(a) = \alpha(a) + \alpha(a')k$.
- $\hat{\alpha}(a'') = \alpha(a'') + \alpha(a')(1-k)$.
- $\hat{\alpha}(\hat{a}) = \alpha(\hat{a})$ for all $\hat{a} \in A \setminus \{a, a', a''\}$.

Notice that $\hat{\alpha}$ is optimal because $\hat{\alpha} \sim_{v_1} \alpha$ and $\hat{\alpha} \sim_{v_2} \alpha$. Moreover, $|S_{\hat{\alpha}}| = |S_\alpha| - 1$. Therefore, if $|S_\alpha| = n > 2$, iterating $n - 2$ times Step 4 we can obtain an optimal action $\tilde{\alpha}$ with $|S_{\tilde{\alpha}}| = 2$.

Proof of Proposition 7

Take any finite set of utilities $\mathcal{W}$ and assume that $\arg\max_{a \in A} u(a) = A$.

Step 1. *If randomization is strictly beneficial, then there is no utility $v \in \mathcal{W}$ such that $P_v = A$.*

Proof of Step 1. Suppose that randomization is strictly beneficial. Then, there exists $\alpha \in \mathcal{A}$ such that $u(\alpha) > u(a)$ for all $a \in A$. By contradiction, suppose that there exists $v \in \mathcal{W}$ such that $P_v = A$. Then, for all $a \in A$, we have

$$\mathrm{U}\left[\mathbb{E}(a,v),v\right] = \mathrm{U}\left[\mathbb{E}(\alpha,v),v\right] \geq u(\alpha) > u(a) = \mathrm{U}\left[\mathbb{E}(a,v),v\right],$$

where the first equality holds because $P_v = A$ and $\arg\max_{a \in A} u(a) = A$, the second weak inequality by definition of $u(\cdot)$, the third strict inequality by assumption and the last equality because $P_v = A$. Therefore, we obtained a contradiction.

Step 2. *If there is no utility $v \in \mathcal{W}$ such that $P_v = A$, then randomization is strictly beneficial.*

Proof of Step 2. Take an action $\alpha \in \mathcal{A}$ with $S_\alpha = A$. Notice that for all $\tilde{a} \in A$, and for all $v \in \mathcal{W}$,

$$\mathrm{U}\left[\mathbb{E}(\alpha,v),v\right] \geq \mathrm{U}\left[\min_{a \in A} \mathbb{E}(a,v),v\right] \geq u(\tilde{a}),$$

where the first weak inequality holds because the function $\mathrm{U}\left[\cdot,v\right]$ is strictly increasing in the first argument and the second weak inequality because $\arg\max_{a \in A} u(a) = A$ and by definition of $u(\cdot)$. Moreover, if the first inequality holds as equality, then the second inequality is strict. Otherwise, we would have $P_v = A$. Therefore, incentives to randomize are strict.

Proof of Proposition 8

Suppose that $\mathcal{W} = \{v_1, v_2\}$, $\arg\max_{a \in A} u(a) = A$ and there is no utility $v \in \mathcal{W}$ such that $P_v = A$.

Step 1. *If action $\alpha \in \mathcal{A}$ is optimal, then $M_\alpha = \{v_1, v_2\}$.*

Proof of Step 1. See Step 1 in the proof of Proposition 6.

Step 2. *If action $\alpha \in \mathcal{A}$ is optimal, then*

$$S_\alpha \subseteq \underset{a \in P_{v_1} \setminus P_{v_2}}{\arg\max}\, \mathrm{U}\left[\mathbb{E}(a, v_2), v_2\right] \cup \underset{a \in P_{v_2} \setminus P_{v_1}}{\arg\max}\, \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right].$$

Proof of Step 2. Suppose that there is an action $a \in S_\alpha$ such that

$$a \notin \underset{a \in P_{v_1} \setminus P_{v_2}}{\arg\max}\, \mathrm{U}\left[\mathbb{E}(a, v_2), v_2\right] \cup \underset{a \in P_{v_2} \setminus P_{v_1}}{\arg\max}\, \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right].$$

We show that action α is not optimal. Assume that $M_\alpha = \{v_1, v_2\}$, otherwise by Step 1 the statement follows. There are three cases:

1. $a \in P_{v_1} \cup P_{v_2}$. Consider a new mixed action $\hat{\alpha}$ such that $\hat{\alpha}(a) = 0$ and for all $a' \neq a$, $\hat{\alpha}(a') = \alpha(a') + \alpha(a)(|A| - 1)^{-1}$. Notice that for all actions $a' \in A$, $a' \succsim_{v_1} a$ and $a' \succsim_{v_2} a$ because $\arg\max_{a \in A} u(a) = A$. Moreover, given that that there is no utility $v \in \mathcal{W}$ such that $P_v = A$, there must be two actions a_1 and a_2 such that $a_1 \succ_{v_1} a$ and $a_2 \succ_{v_2} a$. Therefore, it follows that $\hat{\alpha} \succ_{v_1} \alpha$ and $\hat{\alpha} \succ_{v_2} \alpha$, concluding that α is not optimal.

2. $a \in P_{v_1} \setminus P_{v_2}$. Take an action $a' \in P_{v_1} \setminus P_{v_2}$ such that $a' \succ_{v_2} a$. By assumption, such action exists. Consider a new mixed action $\hat{\alpha}$ such that $\hat{\alpha}(a) = 0$, $\hat{\alpha}(a') = \alpha(a) + \alpha(a')$ and for all $a'' \in A \setminus \{a, a'\}$, $\hat{\alpha}(a') = \alpha(a)$. Given that $a' \succ_{v_2} a$ and $a' \sim_{v_1} a$, it must be that $u(\alpha) = u(\hat{\alpha})$ and $M_{\hat{\alpha}} = \{v_1\}$. By Step 1, there exists another action $\tilde{\alpha}$ such that $u(\tilde{\alpha}) > u(\hat{\alpha}) = u(\alpha)$. Therefore, action α is not optimal.

3. $a \in P_{v_2} \setminus P_{v_1}$. Take an action $a' \in P_{v_2} \setminus P_{v_1}$ such that $a' \succ_{v_1} a$. By assumption, such action exists. Consider a new mixed action $\hat{\alpha}$ such that $\hat{\alpha}(a) = 0$, $\hat{\alpha}(a') = \alpha(a) + \alpha(a')$ and for all $a'' \in A \setminus \{a, a'\}$, $\hat{\alpha}(a') = \alpha(a)$. Given that $a' \succ_{v_1} a$ and $a' \sim_{v_2} a$, it must be that $u(\alpha) = u(\hat{\alpha})$ and $M_{\hat{\alpha}} = \{v_2\}$. By Step 1, there exists another action $\tilde{\alpha}$ such that $u(\tilde{\alpha}) > u(\hat{\alpha}) = u(\alpha)$. Therefore, action α is not optimal.

Step 3. *Action $\alpha \in \mathcal{A}$ is optimal if the following statements are true:*

(1) $M_\alpha = \{v_1, v_2\}$

(2) $S_\alpha \subseteq \arg\max_{a \in P_{v_1} \setminus P_{v_2}} \mathrm{U}\left[\mathbb{E}(a, v_2), v_2\right] \cup \arg\max_{a \in P_{v_2} \setminus P_{v_1}} \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right].$

Proof of Step 3. Consider any other mixed action $\tilde{\alpha}$ and suppose by contradiction that $u(\tilde{\alpha}) > u(\alpha)$. By (1), it must be that $\tilde{\alpha} \succ_{v_1} \alpha$ and $\tilde{\alpha} \succ_{v_2} \alpha$. By (2), $\tilde{\alpha} \succ_{v_1} \alpha$ implies that

$$\sum_{a \in P_{v_2} \setminus P_{v_1}} \tilde{\alpha}(a) > \sum_{a \in P_{v_2} \setminus P_{v_1}} \alpha(a).$$

Similarly, by (2), $\tilde{\alpha} \succ_{v_2} \alpha$ implies that

$$\sum_{a \in P_{v_1} \setminus P_{v_2}} \tilde{\alpha}(a) > \sum_{a \in P_{v_1} \setminus P_{v_2}} \alpha(a).$$

Therefore, it must be that

$$\sum_{a \in P_{v_2} \setminus P_{v_1}} \tilde{\alpha}(a) + \sum_{a \in P_{v_1} \setminus P_{v_2}} \tilde{\alpha}(a) > \sum_{a \in P_{v_2} \setminus P_{v_1}} \alpha(a) + \sum_{a \in P_{v_1} \setminus P_{v_2}} \alpha(a) = 1,$$

where the last equality holds by (2). Consequently, we obtained a contradiction.

Proof of Proposition 9

Assume that $A = \{a, b\}$ and $\mathcal{W} = \{v_1, v_2\}$.

Step 1. *If randomization is strictly beneficial, then there is no utility $v \in \{v_1, v_2\}$ such that $P_v = \{a, b\}$.*

Proof of Step 1. See Step 1 in the proof of Proposition 7.

Step 2. *If randomization is strictly beneficial, then either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$.*

Proof of Step 2. Suppose that it is not true that either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$. If $a \succsim_{v_1} b$ and $a \succsim_{v_2} b$, for all mixed actions $\alpha \in \mathcal{A}$, and for all $v \in \{v_1, v_2\}$, we have

$$\mathrm{U}\left[\mathbb{E}(a, v), v\right] \geq \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha),$$

where the first inequality holds because $a \succsim_{v_1} b$ and $a \succsim_{v_2} b$, the second inequality by definition of $u(\cdot)$. Consequently, $u(a) \geq u(\alpha)$ and randomization is not strictly beneficial. If instead $b \succsim_{v_1} a$ and $b \succsim_{v_2} a$, for all mixed actions $\alpha \in \mathcal{A}$, and for all $v \in \{v_1, v_2\}$, we have

$$\mathrm{U}\left[\mathbb{E}(b, v), v\right] \geq \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha),$$

where the first inequality holds because $b \succsim_{v_1} a$ and $b \succsim_{v_2} a$, the second inequality by definition of $u(\cdot)$. Consequently, $u(b) \geq u(\alpha)$ and randomization is not strictly beneficial.

Step 3. *Randomization is strictly beneficial if the following statements are true:*

(1) *There is no utility $v \in \{v_1, v_2\}$ such that $P_v = \{a, b\}$.*

(2) *Either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$.*

Proof of Step 3. By (1), either $M_a = \{v_1\}$ and $M_b = \{v_2\}$, or $M_a = \{v_2\}$ and $M_b = \{v_1\}$. Without loss of generality, assume $M_a = \{v_1\}$ and $M_b = \{v_2\}$. Then it must be that $b \succ_{v_1} a$. Otherwise, by (2) we have $a \succ_{v_1} b$ and $b \succ_{v_2} a$. This implies that

$$\mathrm{U}\left[\mathbb{E}(b, v_2), v_2\right] > \mathrm{U}\left[\mathbb{E}(a, v_2), v_2\right] > \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right] > \mathrm{U}\left[\mathbb{E}(b, v_1), v_1\right],$$

where the first strict inequality holds because $b \succ_{v_2} a$, the second strict inequality because $M_a = \{v_1\}$, and the third strict inequality because $a \succ_{v_1} b$. However, $\mathrm{U}\left[\mathbb{E}(b, v_2), v_2\right] > \mathrm{U}\left[\mathbb{E}(b, v_1), v_1\right]$ contradicts $M_b = \{v_2\}$. Therefore, it must be that $b \succ_{v_1} a$ and by (2) $a \succ_{v_2} b$. Without loss of generality, assume that $u(a) \geq u(b)$. Define a new mixed action α_λ parametrized by $\lambda \in (0, 1)$ such that $\alpha_\lambda(a) = 1 - \lambda$ and $\alpha_\lambda(b) = \lambda$. Notice that for any $\lambda \in (0, 1)$, $\alpha_\lambda \succ_{v_1} a$. Moreover, for λ small enough, $M_{\tilde{\alpha}} = \{v_1\}$. Therefore,

$$u(\alpha_\lambda) = \mathrm{U}\left[\mathbb{E}(\alpha_\lambda, v_1), v_1\right] > \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right] = u(a) \geq u(b),$$

proving that randomization is strictly beneficial.

Proof of Corollary 1

If incentives to randomize are strict, by Proposition 9 either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$. Given that a is a mean-preserving spread of b, then either v_1 is strictly concave and v_2 is strictly convex, or vice versa. In any case, a C-EU DM is neither risk averse, nor risk seeking.

Proof of Corollary 2

By Proposition 6, if α is optimal, then $M_\alpha = \{v_1, v_2\}$. Suppose that $M_\alpha = \{v_1, v_2\}$. Given that randomization is strictly beneficial, either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$. Without loss of generality, assume $a \succ_{v_1} b$ and $b \succ_{v_2} a$. Consider any other mixed action $\tilde{\alpha} \neq \alpha$. If $\tilde{\alpha}(a) > \alpha(a)$, then $\alpha \succ_{v_2} \tilde{\alpha}$ and $u(\alpha) > u(\tilde{\alpha})$. If instead $\alpha(a) > \tilde{\alpha}(a)$, then $\alpha \succ_{v_1} \tilde{\alpha}$ and $u(\alpha) > u(\tilde{\alpha})$. Therefore α is optimal and unique.

Proof of Corollary 3

Let $\mathcal{W}_A = \{v_1, v_2\}$ and $\mathcal{W}_B = \{w_1, w_2\}$. Assume that (α, β) is a mixed Nash equilibrium of G.

Step 1. *If* $u_B(\beta, \alpha) > \max\{u_B(b_1, \alpha), u_B(b_2, \alpha)\}$, *then* $\alpha \in \bar{A}^{\mathrm{o}}$.

Proof of Step 1. Suppose that $\alpha \notin \bar{A}^{\mathrm{o}}$. Then, either $b_1 \succsim_{w,\alpha} b_2$ for all utilities $w \in \mathcal{W}_B$ or vice versa. In both cases, by Proposition 9, incentives to randomize for player B are not strict.

Step 2. *If* $\alpha \in \bar{A}^{\mathrm{o}}$, *then* $u_B(\beta, \alpha) > \max\{u_B(b_1, \alpha), u_B(b_2, \alpha)\}$.

Proof of Step 2. Suppose that $\alpha \in \bar{A}^{\mathrm{o}}$. Then, either $b_1 \succ_{w_1,\alpha} b_2$ and $b_2 \succ_{w_2,\alpha} b_1$, or $b_2 \succ_{w_1,\alpha} b_1$ and $b_1 \succ_{w_2,\alpha} b_2$. Without loss of generality, assume that the first case holds. By Proposition 9, it is enough to show that there is no utility $w \in \mathcal{W}_B$ such that $P_{w,\alpha} = S_B$. By contradiction, suppose that $P_{w_1,\alpha} = S_B$. Then, we have

$$\begin{aligned} u_B(b_1, \alpha) &= \mathrm{U}\left[\mathbb{E}_\alpha(b_1, w_1), w_1\right] \\ &> \mathrm{U}\left[\mathbb{E}_\alpha(\beta, w_1), w_1\right] \\ &\geq u_B(\beta, \alpha), \end{aligned}$$

where the first equality holds because $P_{w_1,\alpha} = S_B$, the second strict inequality because $b_1 \succ_{w_1,\alpha} b_2$ implies $b_1 \succ_{w_1,\alpha} \beta$, and the third weak inequality by definition of $u_B(\cdot, \alpha)$. However, given that (α, β) is a mixed Nash equilibrium of G, this is not possible.

Step 3. *$\beta \in \bar{B}^{\mathrm{o}}$ if and only if then $u_A(\alpha, \beta) > \max\{u_A(a_1, \beta), u_A(a_2, \beta)\}$.*

Proof of Step 3. It follows from the same arguments that we use in Steps 1 and 2.

Proof of Corollary 4

Let $X \subseteq B^{\mathrm{o}}$ and assume that for all $\beta \in X$, incentives to randomize are strict for player A. By Corollary 2, the unique optimal mixed action α of player A satisfies $M_\alpha = \mathcal{W}_A$. This implies that α satisfies $\mathbb{E}_\beta[\alpha, v_A] = \mathbb{E}_\beta[\alpha, w_A]$. Solving this equation for α yields the desired result.

Proof of Corollary 5

Let $\mathcal{W}_A = \{v_1, v_2\}$ and $\mathcal{W}_B = \{w_1, w_2\}$. Assume that (α, β) is a mixed Nash equilibrium of G.

Step 1. *If (α, β) is strict, then it is efficient.*

Given that (α, β) is a strict mixed Nash equilibrium, given the correct conjectures, incentives to randomize are strict for both players, and actions α and β are optimal. By Corollary 2, α and β are the unique optimal actions and therefore they are efficient. Consequently, the equilibrium (α, β) is efficient.

Step 2. *If (α, β) is is efficient, then it is neither weak nor partially strict provided that $\bar{A}^{\mathrm{o}}$ and $\bar{B}^{\mathrm{o}}$ are non-empty.*

Suppose that (α, β) is a weak or partially strict mixed Nash equilibrium. We show that (α, β) is not efficient. Given that (α, β) is not a strict equilibrium, there exists one player for which randomization is not strictly beneficial in equilibrium. Without loss of generality, assume that this is true for player A. That is, given the correct conjecture β, we have

$$u_A(\alpha, \beta) = \max\{u_A(a_1, \beta), u_A(a_2, \beta)\}.$$

By Corollary 3, it must be that either $\beta = \min(\bar{B})$ or $\beta = \max(\bar{B})$. Given that $\bar{B}^o$ is non-empty, these two quantities are distinct. In both cases, it must be the case that one utility in $\mathcal{W}_A$ is indifferent between the pure actions a_1 and a_2, while the other utility in $\mathcal{W}_A$ strictly prefers one of the two actions. Consequently, one of the two pure actions Pareto dominates action α in $\mathcal{W}_A$, proving that action α is not efficient. Therefore, also the mixed Nash equilibrium (α, β) is not efficient.

References

Agranov, M., Healy, P., and Nielsen, K. (2021). Stable randomization. *SSRN working paper*.

Agranov, M. and Ortoleva, P. (2017). Stochastic choice and preferences for randomization. *Journal of Political Economy*, 125(1):40–68.

Agranov, M. and Ortoleva, P. (2020). Ranges of preferences and randomization. Technical report, mimeo, Princeton Univ.

Agranov, M. and Ortoleva, P. (2022). Revealed preferences for randomization: An overview. *AEA Papers and Proceedings*, 112:426—-430.

Allen, R. and Rehbeck, J. (2021). A generalization of quantal response equilibrium via perturbed utility. *Games*, 12(1).

Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., and Marinacci, M. (2017). Mixed extensions of decision problems under uncertainty. *Economic Theory*, 63(4):827–866.

Börgers, T. (2017). (no) foundations of dominant-strategy mechanisms: a comment on chung and ely. *Review of Economic Design*, 21:73–82.

Calford, E. M. (2021). Mixed strategies and preference for randomization in games with ambiguity averse agents. *Journal of Economic Theory*, 197:105326.

Cerreia-Vioglio, S. (2009). Maxmin expected utility on a subjective state space: Convex preferences under risk. *preprint*.

Cerreia-Vioglio, S., Dillenberger, D., and Ortoleva, P. (2015). Cautious expected utility and the certainty effect. *Econometrica*, 83(2):693–728.

Cerreia-Vioglio, S., Dillenberger, D., and Ortoleva, P. (2020). An explicit representation for disappointment aversion and other betweenness preferences. *Theoretical Economics*, 15(4):1509–1546.

Chung, K. and Ely, J. (2007). Foundations of dominant-strategy mechanisms. *Review of Economics Studies*, 74(2):447–476.

Dillenberger, D. (2010). Preferences for one-shot resolution of uncertainty and allais-type behavior. *Econometrica*, 78(6):1973–2004.

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

Gul, F. (1991). A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society*, pages 667–686.

Luce, R. D. and Raiffa, H. (1957). *Games and decisions: Introduction and critical survey*. Wiley, New York.

Maccheroni, F. (2002). Maxmin under risk. *Economic Theory*, 19(4):823–831.

Machina, M. J. (1985). Stochastic choice functions generated from deterministic preferences over lotteries. *Economic Journal*, 95(379):575—-94.

Marinacci, M. (2000). Ambiguous games. *Games and Economic Behavior*, 31(2):191–219.

Rothschild, M. and Stiglitz, J. E. (1970). Increasing risk: I. a definition. *Journal of Economic theory*, 2(3):225–243.