Rationality and Dominance*

Bruno Salcedo[†]

Winter 2020

This note presents a classic result in decision theory under uncertainty that characterizes the implications of the expected utility hypothesis. Suppose that a researcher knows the actions available to a decision maker, the possible states of the world, and the agents utility function. However, the researcher does not know the agent's subjective beliefs about the state of the world. The researcher can sometimes still make informative predictions about the behavior of the agent by ruling out actions that would never be optimal regardless of the beliefs of the agent. We will discuss ways to characterize the implications of rationality using different notions dominance between actions. Finally, we will briefly discuss how our analysis can help recover both utility and belief parameters from choice data.

The use of dominance to characterize rationality can be traced back to Borel (1921), who used dominance to solve a specific zero-sum game. The proof techniques we use in this note and the first general result were introduced by Abraham Wald in a series of papers starting with Wald (1939). Stein (1955) proved the first general result for zero-sum games. The specific result in the currest document is due to Pearce (1984), who developed it independently of the previous work on the topic. For more details of the early history of the notions of rationality and dominance see Rukhin (1995).

*Lecture notes for Econ 9602 at Western University.

[†]Department of Economics, Western University · brunosalcedo.com · bsalcedo@uwo.ca

1. Rationality and Strict Dominance

We are interested in a class of decision problems under uncertainty described as follows. An economic agent that must choose an *action* from a fixed nonempty set. The *utility* of the agent depends both on the action taken and on the true *state of the world*. The agent is uncertain about the true state of the world. They only know that the it belongs to a fixed nonempty set of possible states of the world. We maintain the assumption that the agent is rational, in that they maximize their expected utility given their beliefs about the state of the world. The main components of the problem can be captured by the following formal definition.

Definition 1.1 A *decision problem* (under uncertainty) is a tupple (A, X, u), where A is a nonempty set of actions, X is a nonempty set of possible states of the world, and $u : A \times X \to \mathbb{R}$ is a Bernoulli utility function.

For simplicity, we assume that both A and X are finite. All the results continue to hold as long as A and X are compact and u is continuous. Suppose that the agent's subjective beliefs about the state are $\beta \in \Delta X$. Being rational, they will choose actions that maximize their expected utility given by

$$U(a, \beta) := \sum_{x \in X} \beta(x) u(a, x).$$

Such actions are called best responses. Formally, an action a is a *best response* to a belief β if and only if $U(a, \beta) \geq U(a', \beta)$ for every other action a. The set of best responses to β is denoted by $\mathrm{BR}(\beta)$.

Note that our description of the environment does not specify the agent's beliefs. That is because we assume that beliefs are subjective, and the researcher cannot observe them directly. Instead, we allow the agent to hold arbitrary beliefs. Despite this fact, rationality can help to make behavior predictions, because there are actions that are not best responses to any beliefs. Rational players would never choose such actions. In general, a rational player can only choose actions that are rational according to the following definition.

Definition 1.2 An action $a \in A$ is *rational* if there exists some belief $\beta \in \Delta X$ such that $a \in \mathrm{BR}(\beta)$.

	x $[p]$	y $[1-p]$
a	1	0
b	0	1
c	γ	γ

Figure 1 – Utility function for decision problem with $A = \{a, b, c\}$ and $X = \{x, y\}$. If $\gamma < 1/2$, then c is no optimal regardless of the agent's beliefs.

Example 1.1 Suppose that $A = \{a, b, c\}$, $X = \{x, y\}$, and the agent's utility is given in Figure 1, where $\gamma < 1/2$. Beliefs are captured by a number $p := \beta(x) \in [0, 1]$. Note that $U(a, \beta) = p \geq 0$, $U(b, \beta) = (1 - p) \geq 0$, and $U(c, \beta) = \gamma < 1/2$. Hence $\mathrm{BR}(\beta) = \{a\}$ if $p > 1/2$, $\mathrm{BR}(\beta) = \{b\}$ if $p < 1/2$, and $\mathrm{BR}(\beta) = \{a, b\}$ if $p = 1/2$. There is no beliefs for which c is optimal. Hence, if the decision maker is rational, they would never choose c.

One issue with Definition 1.2 is that it is difficulty to verify directly whether an action is rational. In principle, one would have to completely characterize $\mathrm{BR}(\cdot)$ for all possible beliefs. In problems with more than a few states, this can be a daunting task. Fortunately, the set of rational actions can be characterized using dominance relations. Informally, one action dominates another if it gives a higher expected utility regardless of the beliefs of the agent. Before defining dominance formally, let us revisit our example.

Example 1.1 (Continued). First, suppose that $\gamma < 0$. In that case, it is easy to see that a results in higher utility than c, regardless of the state. Consequently, a results in a higher expected utility than c, regardless of the agent's beliefs. In such cases we say that *a strictly dominates c*.

Now, suppose that $0 < \gamma < 1/2$. In this case, c is not strictly dominated by a nor b. However, the agent could make a random choice choosing a and b with probability $1/2$ each (e.g., using a fair coin). Doing so would result in the expected utility

$$\frac{1}{2}U(a, \beta) + \frac{1}{2}U(b, \beta) = \frac{1}{2}p + \frac{1}{2}(1 - p) = \frac{1}{2} > \gamma = U(c, \gamma).$$

Therefore c cannot be a best response to any belief.

A *mixed action* is a probability measure $\alpha \in \Delta A$. With slight abuse of notation, we identify each pure action with the degenerated mixed action that assigns full probability to it. Hence, we can think of pure actions as points in ΔA. Moreover, we extend the use of U to denote extended utility from mixed actions as follows

$$U(\alpha, x) = \sum_{a \in A} \alpha(a) u(a, x) \qquad \text{and} \qquad U(\alpha, \beta) = \sum_{a \in A} \sum_{x \in X} \alpha(a) \beta(x) u(a, x).$$

At this point in the course, we do not care whether the agent makes random choices. Mixed actions are used only as a technical tool to determine which pure actions are rational and which re not.

Definition 1.3 Given two pure or mixed actions $\alpha, \alpha' \in \Delta A$, action α is said to *strictly dominate* α' if and only if $U(\alpha, x) > U(\alpha', x)$ for every $x \in X$.

Say that an action is *dominated* if it is strictly dominated by some other action. Otherwise say that it is *undominated*. As it turns out, the set of rational actions corresponds to the set of undominated actions. The importance of the notion of dominance is that it is much more tractable computationally.

Theorem 1.1 *A pure action is rational if and only if it is undominated.*

The following lemmas from will be useful later in the course, when we discuss rationalizability in strategic games.

Lemma 1.2 *If a mixed action assigns positive probability to a dominated pure action, then it is dominated.*

Lemma 1.3 (Dufwenberg and Stegeman (2002)) *Every dominated action is strictly dominated by a mixed action supported on the set of undominated pure actions.*

2. Proofs

The proofs of Theorem 1.1 and Lemma 1.3 require some preliminary definitions. Let $n = \|X\|$ be the number of possible states of the world. We can identify both the actions and the beliefs of the agent with vectors in $\mathbb{R}^n$. For each pure or mixed action α, let $v(\alpha) = (U(\alpha, x))_{x \in X} \in \mathbb{R}^n$ denote the vector of payoffs that can arise from choosing α. For each belief β, let $p(\beta) = (\beta(x))_{x \in X} \in \mathbb{R}^n$ denote the vector of the probabilities of the possible states of the world according to β. Then, we can write expected utility as an inner product $U(\alpha, \beta) = p(\beta)^\mathsf{T} v(\alpha)$.

We also need to define some subsets of $\mathbb{R}^n$. First, let $V = \{v(\alpha) \mid \alpha \in \Delta A\}$ be the set of feasible payoffs. Now, consider an arbitrary pure or mixed action α. Let $W(\alpha) = \{w \in \mathbb{R}^n \mid w \gg v(\alpha)\}$ denote the set of payoff vectors that are coordinate-wise strictly greater than $v(\alpha)$. Finally, let $\overline{W}(\alpha) = \{w \in \mathbb{R}^n \mid w \geq v(\alpha)\}$ be the closure of $W(\alpha)$. Note that α is dominated if and only if $V \cap W(\alpha) \neq \emptyset$. With these preliminaries we can proceed to the actual proofs.

Proof of Theorem 1.1. (Dominated $\Rightarrow$ not rational) Suppose $a \in A$ is dominated by $\alpha \in \Delta A$, and consider any belief $\beta \in \Delta X$. Note that

$$\begin{aligned} U(a, \beta) &= \sum_{x \in X} \beta(x) u(a, x) < \sum_{x \in X} \beta(x) u(a, x) U(\alpha, x) \\ &= \sum_{x \in X} \sum_{a' \in A} \alpha(a') \beta(x) u(a', x) = U(\alpha, \beta). \end{aligned}$$

(Not dominated $\Rightarrow$ rational) Fix an undominated pure action a^*. The sets V and $W(a)$ are convex and disjoint (see Figure 2). Hence, the Minkowski's Separating-Hyperplane Theorem implies that there exists some no-null vector $q \in \mathbb{R}^n \setminus \{0\}$ such that

$$\sup \left\{ q^\mathsf{T} v \mid v \in V \right\} \leq \inf \left\{ q^\mathsf{T} w \mid w \in W(a) \right\} = \inf \left\{ q^\mathsf{T} w \mid w \in \overline{W}(a) \right\}, \tag{1}$$

Note that $v(a^*)$ belongs to $\overline{W}(a)$. Therefore,

$$q^\mathsf{T} v(a^*) \geq \inf \left\{ q^\mathsf{T} w \mid w \in W(a) \right\}. \tag{2}$$

Conditions (1) and (2) together imply that $q^\mathsf{T} v(a^*) \geq q^\mathsf{T} v(a)$ for every $a \in A$. The rest of the proof shows that there exists a number $\eta > 0$ such that $p := \eta \cdot q$ is a

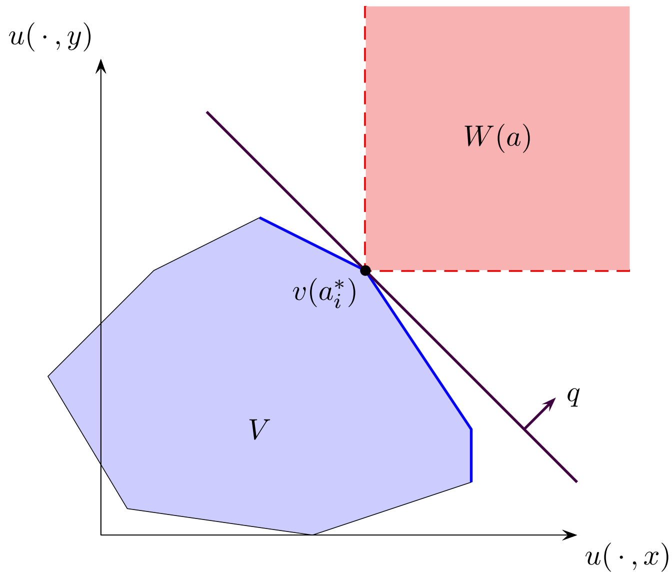

Figure 2 – Example with $X = \{x, y\}$ illustrating Theorem 1.1. An action is rational if and only if it is undominated (its payoff vector in in the North-East boundary of V). The undominated action a_i^* is a best response to a belief parallel to q.

probability vector. Consequently, there exists some $\beta \in \Delta X$ such that $p(\beta) = p$ and therefore

$$U(a, \beta) = (\eta \cdot q)^\intercal v(a) \leq (\eta \cdot q)^\intercal v(a^*) = U(a^*, \beta),$$

for every $a \in A$. That is, a^* is a best response to β.

For each coordinate $j = 1, \ldots, n$, let $w^j \in \mathbb{R}^n$ be the vector with $w^j_j = 1$ and $w^j_k = 0$ for $k \neq j$. Note that $v(a^*) + w^j \in \overline{W}(a)$. Being that $v(a^*) \in V$, condition (1) implies that

$$q^\intercal v(a^*) \leq q^\intercal (v(a^*) + w^j) = q^\intercal v(a^*) + q_j.$$

Therefore, we must have $q_j \geq 0$ for all $j = 1, \ldots, n$. Moreover, since $q \neq 0$, we must have $\sum_{j=1}^n q_j > 0$. Hence, we can simply set $\eta = 1/(\sum_{j=1}^n q_j)$ so that $\eta \cdot q$ belongs to the n-dimensional simplex, thus completing the proof. ∎

Proof of Lemma 1.2. Consider any pure action a^0 and any mixed action α such that $\alpha(a^0) > 0$. Suppose that a^0 is strictly dominated by a mixed action $\alpha^* \in \Delta A$. We will show that α is also strictly dominated by the mixed action α' given by

$\alpha'(a^0) = \alpha(a^0)\alpha^*(a^0)$ and

$$\alpha'(a) = \alpha(a) + \alpha(a^0)\alpha^*(a),$$

for $a \neq a^0$. First, we have to show that α' is indeed a mixed action. Clearly, $\alpha'(a) \geq 0$ for every action a. Moreover,

$$\sum_{a\in A} \alpha'(a) = \sum_{a\neq a^0} \alpha(a) + \alpha(a^0)\cancelto{1}{\sum_{a\in A} \alpha^*(a)} = \sum_{a\in A} \alpha(a) = 1.$$

Hence, $\alpha' \in \Delta A$. Now, note that for every state $x \in X$,

$$\begin{aligned} U(\alpha', x) &= \sum_{a\neq a^0} \alpha(a)u(a,x) + \alpha(a^0) \sum_{a\in A} \alpha^*(a)u(a,x) \\ &= \sum_{a\neq a^0} \alpha(a)u(a,x) + \alpha(a^0)U(\alpha^*, x) \\ &> \sum_{a\neq a^0} \alpha(a)u(a,x) + \alpha(a^0)u(a^0, x) = U(\alpha, x). \end{aligned}$$

Hence, α is strictly dominated by α'. ■

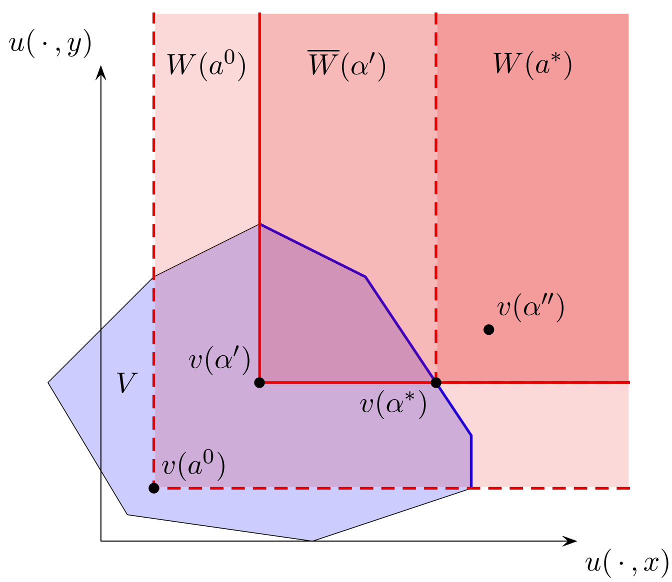

Figure 3 – Example with $X = \{x, y\}$ illustrating Lemma 1.3. Since a^0 is strictly dominated by α', it is also dominated by some mixed action α^* that assigns full probability to undominated actions.

Proof of Lemma 1.3. Suppose that a given action a^0 is strictly dominated by α'. Fix an arbitrary state x and a pure or mixed strategy α^* that solves the program

$$\max_{\alpha \in \Delta A} U(\alpha, x) \qquad \text{s.t.} \quad v(\alpha) \in V \cap \overline{W}(\alpha'). \tag{3}$$

See Figure 3. Since V is compact and $\bar{W}(\alpha')$ is closed and nonempty, $V \cap \bar{W}(\alpha')$ is compact and thus α^* is well defined. Suppose towards a contradiction that α^* is strictly dominated by some $\alpha'' \in \Delta A$. We would have $v(\alpha'') \in W(\alpha^*) \subseteq V \cap \overline{W}(\alpha')$, and $U(\alpha'', x) > U(\alpha^*, x)$. But this would contradict the fact that α^* is a solution to (3) ▼. Hence, α^* is undominated, and the result thus follows from Lemma 1.2. ■

References

Borel, E. (1953/1921). The theory of play and integral equations with skew symmetric kernels [La théorie du jeu et les équations intégrales à noyausymétrique gauche]. *Econometrica*, 21(1):97–100. Translated from French by Savage, L. J.

Dufwenberg, M. and Stegeman, M. (2002). Existence and uniqueness of maximal reductions under iterated strict dominance. *Econometrica*, 70(5):2007–2023.

Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.

Rukhin, A. L. (1995). Admissibility: Survey of a concept in progress. *International Statistical Review/Revue Internationale de Statistique*, pages 95–115.

Stein, C. (1955). A necessary and sufficient condition for admissibility. *The Annals of Mathematical Statistics*, 26(3):518–522.

Wald, A. (1939). Contributions to the theory of statistical estimation and testing hypotheses. *The Annals of Mathematical Statistics*, 10(4):299–326.