Dominance and Optimality

by Xienan Cheng* ⓡ Tilman Börgers†

June 27, 2025

Online Appendix

A. Value Functions for Experiments

Proposition 1. *A function $v : \Gamma_{\mathcal{S}} \to \mathbb{R}$ is finitely convex if and only if there exists a decision problem (A, u) such that v is the value function for this decision problem.*

Proof. The "if" direction of this claim is standard. We omit the proof. We prove the "only if" direction. That is, we prove that if v is finitely convex, then there exists a decision problem (A, u) such that v is the value function for this decision problem. The proof describes how to construct such a decision problem. Define $\mathcal{E}$ to be the convex hull of the epigraph of v:

$$\mathcal{E} \equiv co\{(\gamma, x) | \gamma \in \Gamma_{\mathcal{S}} \wedge x \geq v(\gamma)\}.$$

The set $\mathcal{E}$ is finitely generated in the sense of Rockafellar (1970, p. 170) because it is the convex hull of the points $(\gamma, v(\gamma))$ for $\gamma \in \Gamma_{\mathcal{S}}$, and of the direction $(\vec{0}, 1)$ where $\vec{0}$ denotes the zero vector in $\mathbb{R}^{|\Omega|}$. By Theorem 19.1 in Rockafellar (1970, p. 171) the set $\mathcal{E}$ is polyhedral, which means that it equals the set of solutions of a finite system of inequalities of the form:

$$a^T \gamma + bx \leq c,$$

where $a \in \mathbb{R}^{|\Omega|}$ and $b, c \in \mathbb{R}$.

Note that every such inequality must satisfy $b \leq 0$. This is because $\mathcal{E}$ is not bounded from above in its last component, and therefore, if $b > 0$, we could always find an element of $\mathcal{E}$ where the last component is so large that the inequality is violated.

*Guanghua School of Management, Peking University, xienancheng@gsm.pku.edu.cn.
†Department of Economics, University of Michigan, tborgers@umich.edu.

We can thus distinguish between inequalities for which $b = 0$ and inequalities for which $b < 0$. Consider inequalities for which $b = 0$. They are of the form:

$$a^T\gamma \leq c.$$

These inequalities must be satisfied for all $\gamma \in co(\Gamma_{\mathcal{S}})$. Thus, we can drop these inequalities, and instead describe the set $\mathcal{E}$ as the set of all pairs (γ, x) such that $\gamma \in co(\Gamma_{\mathcal{S}})$, and such that a finite set of inequalities of the form

$$a^T\gamma + bx \leq c,$$

hold, where now for each such inequality we have: $b < 0$.

Now consider any point $(\bar{\gamma}, v(\bar{\gamma}))$ where $\bar{\gamma} \in \Gamma_{\mathcal{S}}$. Clearly, this point is an element of $\mathcal{E}$, and therefore satisfies all inequalities that describe $\mathcal{E}$. We now claim that this point satisfies at least one inequality with equality. Suppose not. Then there would be an $x \in \mathbb{R}$ with $(\bar{\gamma}, x) \in \mathcal{E}$ and $x < v(\bar{\gamma})$. But this contradicts the definition of $\mathcal{E}$ together with the finite convexity of v. Any element $(\bar{\gamma}, x)$ of $\mathcal{E}$ can be written as a convex combination of the form:

$$(\bar{\gamma}, x) = \sum_{\gamma \in \Gamma_{\mathcal{S}}} \lambda_\gamma (\gamma, x_\gamma)$$

where $\lambda_\gamma \geq 0$ for all $\gamma \in \Gamma_{\mathcal{S}}$, $\sum_{\gamma \in \Gamma_{\mathcal{S}}} \lambda_\gamma = 1$, and $x_\gamma \geq v(\gamma)$ for all $\gamma \in \Gamma_{\mathcal{S}}$. We thus have:

$$\bar{\gamma} = \sum_{\gamma \in \Gamma_{\mathcal{S}}} \lambda_\gamma \gamma,$$

and

$$x \geq \sum_{\gamma \in \Gamma_{\mathcal{S}}} \lambda_\gamma v(\gamma).$$

But the finite convexity of v now implies that the right hand side of this expression is not less than $v(\bar{\gamma})$, and thus we have: $x \geq v(\bar{\gamma})$, which contradicts our initial assumption.

Now consider any point $(\bar{\gamma}, v(\bar{\gamma}))$ where $\bar{\gamma} \in \Gamma_{\mathcal{S}}$. Clearly, this point is an element of $\mathcal{E}$, and therefore satisfies all inequalities that describe $\mathcal{E}$. We now claim that this point satisfies at least one inequality with equality. Suppose not. Then there would be an $x \in \mathbb{R}$ with $(\bar{\gamma}, x) \in \mathcal{E}$ and $x < v(\bar{\gamma})$. But this contradicts the definition of $\mathcal{E}$ together with the finite convexity of v. Any element $(\bar{\gamma}, x)$ of $\mathcal{E}$ can be written as a convex combination of the form:

$$(\bar{\gamma}, x) = \sum_{\gamma \in \Gamma_{\mathcal{S}}} \lambda_\gamma (\gamma, x_\gamma)$$

where $\lambda_\gamma \geq 0$ for all $\gamma \in \Gamma_{\mathcal{S}}$, $\sum_{\gamma\in\Gamma_{\mathcal{S}}} \lambda_\gamma = 1$, and $x_\gamma \geq v(\gamma)$ for all $\gamma \in \Gamma_{\mathcal{S}}$. We thus have:

$$\bar{\gamma} = \sum_{\gamma\in\Gamma_{\mathcal{S}}} \lambda_\gamma \gamma,$$

and

$$x \geq \sum_{\gamma\in\Gamma_{\mathcal{S}}} \lambda_\gamma v(\gamma).$$

But the finite convexity of v now implies that the right hand side of this expression is not less than $v(\bar{\gamma})$, and thus we have: $x \geq v(\bar{\gamma})$, which contradicts our initial assumption.

We thus conclude that for every $\gamma \in \Gamma_{\mathcal{S}}$ one of the inequalities that describe $\mathcal{E}$ holds as equality:

$$\begin{aligned} a^T\gamma + bv(\gamma) &= c \Leftrightarrow \\ -\frac{a^T}{b}\gamma + \frac{c}{b} &= v(\gamma) \end{aligned}$$

Denote by $\frac{\vec{c}}{b}$ the vector that consists of $|\Omega|$ repetitions of $\frac{c}{b}$. Then, because the components of γ add up to 1, we can write this as:

$$\left(-\frac{a}{b} + \frac{\vec{c}}{b}\right)^T \gamma = v(\gamma)$$

Thus, if the decision maker with beliefs γ chooses an action where the utility in each state in Ω is given by the corresponding entry in $-\frac{a}{b} + \frac{\vec{c}}{b}$, then this decision maker's expected utility is payoff is $v(\gamma)$.

We can repeat this construction for every $\gamma \in \Gamma_{\mathcal{S}}$. For each γ we obtain a corresponding action where the vector of utilities corresponding to these actions in each state is given by:

$$-\frac{a(\gamma)}{b(\gamma)} + \frac{\vec{c}(\gamma)}{b(\gamma)},$$

where $a(\gamma)$, $b(\gamma)$ and $c(\gamma)$ are the coefficients of the inequality that γ and $v(\gamma)$ satisfy as equalities. We thus obtain a finite decision problem.

To prove the claim we now claim that for each $\gamma \in \Gamma_{\mathcal{S}}$ the decision maker maximizes expected utility by choosing the action that we have constructed that corresponds to γ. Once we have proven this claim, we may conclude that v is the value function corresponding to this decision problem. But because every element of $\mathcal{E}$ satisfies all the inequalities that describe $\mathcal{E}$, we have for all $\gamma \in \Gamma_{\mathcal{S}}$, and for all relevant inequalities:

$$\begin{aligned} a^T\gamma + bv(\gamma) &\leq c \Leftrightarrow \\ -\frac{a^T}{b}\gamma + \frac{c}{b} &\leq v(\gamma) \end{aligned}$$

This last inequality holds for all actions included in the finite decision problem. This implies that choosing an action that yields $v(\gamma)$ is indeed expected utility maximizing if the decision maker's beliefs are given by γ. □

Proposition 2. *A function $v : \Gamma_{\mathcal{S}} \to \mathbb{R}$ is strictly finitely convex if and only if there exists a decision problem (A, u) such that v is the value function for this decision problem, and for each $\gamma \in \Gamma_{\mathcal{S}}$, there is a unique and distinct optimal action.*

Proof. The "if" direction of this claim is standard. We omit the proof. We prove the "only if" direction. The proof describes how to construct such a decision problem. Define the set $\mathcal{E}$ as in the proof of Proposition 1. The following observation will be crucial:

Lemma 1. *If v is strictly finitely convex, then every $(\gamma, v(\gamma))$ where $\gamma \in \Gamma_{\mathcal{S}}$ is an extreme point of $\mathcal{E}$.*

Proof. Suppose one $(\bar{\gamma}, v(\bar{\gamma}))$ where $\bar{\gamma} \in \Gamma_{\mathcal{S}}$ is not an extreme point. Then there exist distinct $(\gamma_i, x_i) \in \mathcal{E}$ and weights $\lambda_i > 0$, $i = 1, \cdots, n$ such that $n \geq 2$, $\sum_{i=1}^n \lambda_i = 1$, and:

$$(\bar{\gamma}, v(\bar{\gamma})) = \sum_{i=1}^{n} \lambda_i (\gamma_i, x_i).$$

Observe that it is without loss of generality to assume that $\gamma_i \in \Gamma_{\mathcal{S}}$ for all i. If $\gamma_i \notin \Gamma_{\mathcal{S}}$ then we can write (γ_i, x_i) as a convex combination of pairs (γ, x) where $\gamma \in \Gamma_{\mathcal{S}}$ for each pair. This is because $(\gamma_i, x_i) \in \mathcal{E}$ and $\mathcal{E}$ is the convex hull of the set of pairs (γ, x) where $\gamma \in \Gamma_{\mathcal{S}}$.

We now distinguish two cases. The first case is that $\gamma_i = \bar{\gamma}$ for all i. The proof of Proposition 1 shows that we must have $x_i \geq v(\bar{\gamma})$ for all i. (γ_i, x_i) being distinct implies moreover that there exists at least one i such that $x_i > v(\bar{\gamma})$. This implies: $\sum_{i=1}^n \lambda_i x_i > v(\bar{\gamma})$, a contradiction.

The second case is that there exists at least one i such that $\gamma_i \neq \bar{\gamma}$. Then:

$$\sum_{i=1}^{n} \lambda_i x_i \geq \sum_{i=1}^{n} \lambda_i v(\gamma_i) > v(\bar{\gamma}),$$

which is again a contradiction. Here, the first inequality follows as in the first case from the proof of Proposition 1, and the second inequality follows from strict finite convexity of v. □

According to Theorem 2.3 in Bertsimas and Tsitsiklis (2008), $(\bar{\gamma}, v(\bar{\gamma}))$ being an extreme point of the polyhedron $\mathcal{E}$ implies that there exists a supporting hyperplane whose intersection with $\mathcal{E}$ is $\{(\bar{\gamma}, v(\bar{\gamma}))\}$. That is, for every $(\bar{\gamma}, v(\bar{\gamma}))$ with $\bar{\gamma} \in \Gamma_{\mathcal{S}}$, there exists $a(\bar{\gamma}) \in \mathbb{R}^{|\Omega|}$ and $b(\bar{\gamma}), c(\bar{\gamma}) \in \mathbb{R}$ such that

$$a(\bar{\gamma})^T \gamma + b(\bar{\gamma}) x < c(\bar{\gamma})$$

for all $(\gamma, x) \in \mathcal{E} \setminus \{(\bar{\gamma}, v(\bar{\gamma}))\}$, and

$$a(\bar{\gamma})^T \bar{\gamma} + b(\bar{\gamma}) v(\bar{\gamma}) = c(\bar{\gamma}).$$

Note that such inequality must satisfy $b(\bar{\gamma}) < 0$. This is because $\mathcal{E}$ is not bounded from above in its last component, and therefore, if $b(\bar{\gamma}) > 0$, we could always find an element of $\mathcal{E}$ where the last component is so large that the inequality is violated. Furthermore, if $b(\bar{\gamma}) = 0$, then all $(\bar{\gamma}, x)$ in $\mathcal{E}$ make the inequality binding, which leads to a contradiction. So we have $b(\bar{\gamma}) < 0$. Note that the above conditions imply that a different hyperplane corresponds to every extreme point.

Similarly to the construction in the proof of Proposition 1 we now construct the decision problem that has as many actions as there are elements γ of $\Gamma_{\mathcal{S}}$, and where the vector of payoffs for the action corresponding to $\gamma \in \Gamma_{\mathcal{S}}$ is:

$$-\frac{a(\gamma)}{b(\gamma)} + \frac{\vec{c}(\gamma)}{b(\gamma)}$$

By construction then:

$$-\frac{a(\gamma)^T}{b(\gamma)}\gamma + \frac{c(\gamma)}{b(\gamma)} = v(\gamma).$$

for every $\gamma \in \Gamma_{\mathcal{S}}$, and, if $\gamma' \in \Gamma_{\mathcal{S}} \neq \gamma$ then:

$$-\frac{a(\gamma)^T}{b(\gamma)}\gamma' + \frac{c(\gamma)}{b(\gamma)} < v(\gamma').$$

As a result, for every $\gamma \in \Gamma_{\mathcal{S}}$, the action corresponding to γ yields expected utility $v(\gamma)$ and is the only utility maximizing action among all available actions. □

B. Fishburn's Separating Hyperplane Theorem

As we note in the main text, Separating Hyperplane Theorem 2 in the proof of Theorem 2 is Lemma 5 in Fishburn (1975). In Fishburn (1975) Lemma 5 was not explicitly proven. Instead, Fishburn referred the reader to a similar proof in an earlier paper, the proof of Lemma 5 in Fishburn (1974). In the following we explain Fishburn's proof using language and notation that does not refer to the specific application that Fishburn was considering.

Theorem 1 (Separating Hyperplane Theorem 2:)**.** *Let $C \subseteq \mathbb{R}^n$ be non-empty and convex and suppose $C \cap \mathbb{R}^n_- = \emptyset$. Then there exists $\lambda \in \left\{\lambda \in \mathbb{R}^n_+ | \sum_{i=1}^n \lambda_i = 1\right\}$ such that $\lambda \cdot c \geq 0$ for all $c \in C$ and $\lambda \cdot c > 0$ for at least one $c \in C$.*

This theorem would an easy implication of the textbook separating hyperplane theorem due to Minkowski (Ok, 2007, p. 483) were it not for the assertion that $\lambda \cdot c > 0$ for at least one $c \in C$. The following proof shows why this assertion is true.

Proof. Define $\Delta \equiv \left\{\lambda \in \mathbb{R}^n_+ | \sum_{i=1}^n \lambda_i = 1\right\}$. We prove the contrapositive: if for all $\lambda \in \Delta$ we either have: $\lambda \cdot c < 0$ for some $c \in C$ or $\lambda \cdot c = 0$ for all $c \in C$, then $C \cap \mathbb{R}^n_- \neq \emptyset$.

There are two cases in which the assumptions of the contrapositive are satisfied. The first case is that for all $\lambda \in \Delta$ we have: $\lambda \cdot c < 0$ for some $c \in C$. The second case is that there is at least one $\bar{\lambda} \in \Delta$ such that $\bar{\lambda} \cdot c = 0$ for all $c \in C$, and, for all $\lambda \in \Delta$ for which this does not hold, $\lambda \cdot c < 0$ for some $c \in C$.

In the first case the claim follows from a standard separating hyperplane theorem. If there exists no $\lambda \in \Delta$ such that $\lambda \cdot c \geq 0$ for all $c \in C$, then we cannot have $C \cap \mathbb{R}^n_- = \emptyset$. This would contract the Minkowski separating hyperplane theorem (Ok, 2007, p. 483) applied to the case that one of the sets is $\mathbb{R}^n_-$.

We focus on the second case. We prove the claim by induction over n, the dimension of the Euclidean space that we are considering. The claim is trivial if $n = 1$. In this case Δ consists of the single vector $\lambda = 1$, and $\lambda \cdot c = 0$ implies that we must have $c = 0$. Thus, obviously, $c \in \mathbb{R}_-$. Now suppose we had proved the claim for all dimensions $1, 2, \cdots, n-1$. We want to prove it for dimension n.

Pick some $\bar{\lambda} \in \Delta$ such that $\bar{\lambda} \cdot c = 0$ for all $c \in C$. Some components of $\bar{\lambda}$ may be zero. Without loss of generality we assume that, if there are such components, they are the last components of $\bar{\lambda}$, i.e. that we can write $\bar{\lambda}$ as follows:

$$\bar{\lambda} = (\bar{\lambda}_1, \bar{\lambda}_2, \ldots, \bar{\lambda}_m, 0, 0, \ldots, 0) \quad \text{where } 1 \leq m \leq n.$$

Consider the set S of all elements of $\mathbb{R}^n$ such that the first $m-1$ components are zero, all components starting from component $m+1$ are non-positive, and component m can have arbitrary sign. Formally:

$$S \equiv \{c \in \mathbb{R}^n | c_i = 0 \text{ for } i = 1, 2, \ldots, m-1 \text{ and } c_i \leq 0 \text{ for } i = m+1, 2, \ldots, n\}.$$

We will show that:

$$C \cap S \neq \emptyset.$$

First we observe that this claim in fact implies what we have to show: $C \cap \mathbb{R}^n_- \neq \emptyset$. This is because $c \in S$ implies $\bar{\lambda} \cdot c = \bar{\lambda}_m c_m$. By assumption $\bar{\lambda} \cdot c = 0$, and $\bar{\lambda}_m \neq 0$, and thus $c_m = 0$ follows, and therefore $c \in \mathbb{R}^n_-$. Hence it suffices to prove that $C \cap S \neq \emptyset$.

For all $\eta = 0, 1, \ldots m-1$, and for all $\kappa = \eta+1, \eta+2, \ldots, m$ define S^η_κ to be the set of all $c \in \mathbb{R}^n$ such that the first η components of c are zero, all remaining components, except the κ-th component are non-positive, and there is no constraint on the κ-th component. Formally:

$$S^\eta_\kappa \equiv \{c \in \mathbb{R}^n | c_i = 0 \text{ for } i = 1, \cdots, \eta \text{ and } c_i \leq 0 \text{ for } i = \eta+1, \ldots, n \text{ with } i \neq \kappa\}.$$

Obviously, $S = S^{m-1}_m$. We shall prove the assertion by showing that $C \cap S^\eta_\kappa \neq \emptyset$ for all $\eta = 0, 1, \ldots m-1$, and for all $\kappa = \eta+1, \eta+2, \ldots, m$.

		κ						
		1	2	3	4	5	6	7
	0	✓	✓	✓	✓	✓	✓	✓
	1		✓	✓	✓	✓	✓	✓
	2			✓	✓	✓	✓	✓
η	3				✓	✓	✓	✓
	4					✓	✓	✓
	5						✓	✓
	6							✓

TABLE 1. The combinations of η and κ for which the set S^η_κ has been defined

To begin with we visualize in a table the combinations of η and κ for which the sets S_κ^η have been defined. This is done in Table 1. The table is for the case $m = 7$. The rows indicate the value of η, i.e. the number of initial entries of the vectors in S that have to be zero. The columns indicate which entry κ among the remaining entires is allowed to be positive. Checkmarks indicate that the set S_κ^η is well defined.

The inductive assumption of our proof implies that the claim is true for all entries that correspond to the first row in Table 1, i.e. that $C \cap S_\kappa^0 \neq \emptyset$ for all $\kappa = 1, 2, \ldots, m$. Indeed, the stronger claim is true: the inductive assumption implies that for any $\kappa = 1, 2, \ldots, n$ there exists an element c of C such that all components of c other than possibly the κ-th component are non-positive. To see this suppose that we drop the κ-th component from all vectors in C so that we obtain a subset of $\mathbb{R}^{n-1}$. Because C satisfies the assumptions of the contrapositive of this theorem, this new set satisfies the assumptions, too. Therefore, it has non-empty intersection with $\mathbb{R}_-^{n-1}$. Take any element of this intersection, and insert back the κ-th component. Then we have an element of $C \cap S_\kappa^0$.

The proof now shows that if the claim is true for all entries in one row η, then it is also true for all entries in the row $\eta + 1$ in the table. We demonstrate the argument by an example. Suppose in the case illustrated in Table 1 we wanted to prove that the claim holds for all entries in row 3 having proved it for all entries in rows 0, 1, and 2. As an example, let us show that $C \cap S_6^3$ is non-empty. We are going to construct an element of $C \cap S_6^3$. Pick any $r \in C \cap S_3^2$ and $t \in C \cap S_6^2$. Our argument will be that there is a convex combination of r and t that is in S_6^3. Because C is convex, this will be sufficient to prove the claim.

Now r is of the form $(0, 0, r_3, r_4, r_5, r_6, r_7, \ldots, r_n)$ where all entries except r_3 are non-positive. $\bar{\lambda} \cdot r = 0$ implies $r_3 \geq 0$. t is of the form $(0, 0, t_3, t_4, t_5, t_6, t_7, \ldots, t_n)$ where all entries except t_6 are non-positive. $r_3 \geq 0$ and $t_3 \leq 0$ implies there exists a convex combination of r and t, say h, such that $h_3 = 0$. Moreover, the first two components of h are also obviously zero, and all remaining components, except h_6 must be non-positive. Therefore, $h \in S_6^3$.

By iterating this argument, we can conclude that $C \cap S_m^{m-1} \neq \emptyset$ which completes the proof.

□

References

Bertsimas, D. and J. Tsitsiklis (2008). *Introduction to Linear Optimization*. AthenaScientific.

Fishburn, P. C. (1974). Convex stochastic dominance with continuous distribution functions. *Journal of Economic Theory 7*, 143–158.

Fishburn, P. C. (1975). Separation theorems and expected utilities. *Journal of Economic Theory 11*(1), 16–34.

Ok, E. (2007). *Real Analysis with Economic Applications*. Princeton University Press.

Rockafellar, T. (1970). *Convex Analysis*. Princeton: Princeton University Press.