

TOPICS IN GAME THEORY

TAUGHT BY ERAN SHMAYA
NORTHWESTERN UNIVERSITY, FALL 2016

Ludvig Sinander
Northwestern University

This version: 14 November 2018

These notes are based on a theory course for second-year PhD students taught by Eran Shmaya at Northwestern in fall 2016. The topics are mechanism/auction design, supermodular games, and learning.

I thank Eran for teaching a great class and for agreeing to let me share these notes, and Matteo Escudé, Egor Starkov, Ben Vatter, Klavdia Zemlianova and Andrey Zhukov for reporting errors.

Copyright © 2020 Carl Martin Ludvig Sinander.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled ‘GNU Free Documentation License’.

This is a ‘copyleft’ licence. Visit gnu.org/licenses/copyleft to learn more.

Contents

1	Mechanism design	5
1.1	Single agent, one-dimensional types	6
1.1.1	Convex analysis on one-dimensional domains	6
1.1.2	Environment	7
1.1.3	Characterisation of IC	8
1.1.4	Optimal mechanism	10
1.2	Several agents, one-dimensional independent types	12
1.2.1	Environment	12
1.2.2	Bayesian implementation	12
1.2.3	The Myerson–Satterthwaite theorem	14
1.2.4	Strategy-proof implementation	15
1.2.5	Equilibrium via revenue equivalence	17
1.2.6	Optimal mechanism	18
1.2.7	The Bulow–Klemperer theorem	22
1.2.8	Ironing	25
1.2.9	Equivalence of Bayesian and strategy-proof implementation	29
1.3	Single agent, multidimensional types	33
1.3.1	Convex analysis on multidimensional domains	33
1.3.2	Environment	38
1.3.3	Characterisation of IC	39
1.3.4	Optimal mechanism	40
1.4	Several agents, multidimensional types	42
1.4.1	Environment	42
1.4.2	Strategy-proof implementation	42
1.4.3	Roberts’s theorem	43
1.5	Several agents, correlated types	45
1.5.1	Environment	45
1.5.2	Characterisation of BIC (anything goes)	45
1.5.3	Eliciting beliefs	47
1.5.4	More on the ‘anything goes’ result	48
2	Supermodular games	50
2.1	Lattices	50
2.2	Tarski’s fixed-point theorem	55
2.3	Topkis’s monotonicity theorem	60
2.4	Pure-strategy Nash equilibria of supermodular games	62
2.5	Monotone equilibria of supermodular Bayesian games	65

3 Learning	69
3.1 Repeated games with one-sided incomplete information . . .	69
3.1.1 Environment	70
3.1.2 The theorem	70
3.1.3 The lower bound on v_n	74
3.1.4 The upper bound on v_n	77
3.2 Reputation and learning to play Bayesian games	79
References	81

1 Mechanism design

This section doesn't follow any book, but the treatment shares features with Vohra (2011).

We begin with mechanism design in a simple environment: a single indivisible object has to be allocated to (at most) one of n agents whose valuations for it are independently distributed private information. We will sweat this framework pretty hard, and then examine what can be said for more general type spaces and for the case of correlated values.

As a simple motivation, consider a principal allocating a single, indivisible object between two agents. An agent's i 's type is her valuation $t_i \in \mathcal{T}_i \subseteq \mathbf{R}$ for the object. t_i is unknown to the principal, who views it as a random variable T_i distributed with CDF F_i . Further, the principal and agents believe that T_1 and T_2 are distributed independently. Type t_i of agent i derives zero value from not obtaining the object, value t_i from obtaining the object, and maximises her expected utility.

For concreteness, let $T_i \sim \mathcal{U}[0, 1]$. One way to allocate the object is using an auction, e.g. a first-price sealed-bid auction or a second-price sealed-bid auction. As we'll see, these two auctions raise the same revenue, and adding a reserve price may increase revenue further.

In the second-price auction, it's weakly dominant for each player to bid her value, so that 'truthful' bidding constitutes a Bayes–Nash equilibrium. Expected revenue is then $\mathbf{E}(T_1 \wedge T_2) = 1/3$.

In the first-price auction, there is a Nash equilibrium in which both players bid $\alpha(t_i) = t_i/2$ for each $t_i \in [0, 1]$. To verify this, suppose that player j plays α . If type t_i of player i bids a_i , her payoff is

$$\mathbf{P}(\alpha(T_j) < a_i) \cdot (t_i - a_i) = \mathbf{P}(T_j < 2a_i) \cdot (t_i - a_i) = 2a_i(t_i - a_i),$$

which is maximised at $a_i = t_i/2$. Revenue is then $\mathbf{E}(T_1 \vee T_2)/2 = 1/3$, the same.

In an auction with a reserve price \underline{b} , the object is not allocated to either player if they both bid below \underline{b} . In general, this means that the object remains unallocated with positive probability. We shall see that the revenue-maximising way of selling the object is to run an auction with a strictly positive reserve price.

1.1 Single agent, one-dimensional types

1.1.1 Convex analysis on one-dimensional domains

We're going to need some tools from convex analysis. We'll provide results for convex functions defined on intervals of \mathbf{R} ; extensions to multidimensional domains will be covered in section 1.3.1. Proofs, details and related results can be found in Rockafellar (1970).

Definition 1. $f : [a, b] \rightarrow \mathbf{R}$ is convex iff for any $x, x' \in [a, b]$ and $\lambda \in [0, 1]$,

$$f(\lambda x + (1 - \lambda)x') \leq \lambda f(x) + (1 - \lambda)f(x').$$

Proposition 1. Let $\{f_\alpha\}$ be a family of convex functions $[a, b] \rightarrow \mathbf{R}$. Define $f : [a, b] \rightarrow \mathbf{R}$ by $f(x) := \sup_\alpha f_\alpha(x) \forall x \in [a, b]$. Then f is convex.

For $f : [a, b] \rightarrow \mathbf{R}$ convex, $x^* \in \mathbf{R}$ is a subgradient of f at $x \in [a, b]$ iff

$$f(x') \geq f(x) + x^*(x' - x) \quad \forall x' \in [a, b].$$

This means precisely that the affine map $x' \mapsto f(x) + x^*(x' - x)$ is everywhere weakly below f , and equal to f at x . The subdifferential of a convex function $f : [a, b] \rightarrow \mathbf{R}$ is the correspondence $\partial f : [a, b] \rightrightarrows \mathbf{R}$ such that $\partial f(x)$ is the set of all subgradients of f at x . For $f, g \in \mathbf{R}^{[a, b]}$, we write ' $g \in \partial f$ ' as shorthand for ' $g(x) \in \partial f(x)$ for each $x \in [a, b]$ '.

Proposition 2. Let $f : [a, b] \rightarrow \mathbf{R}$ be convex. Then

- (1) $\partial f(x)$ is nonempty, compact and convex for every $x \in [a, b]$.
- (2) $f'_+(x)$ exists at every $x \in [a, b)$, and $f'_-(x)$ exists at every $x \in (a, b]$.¹
- (3) $\partial f(x) = [f'_-(x), f'_+(x)]$ for every $x \in (a, b)$.

Proposition 3. Let $f : [a, b] \rightarrow \mathbf{R}$ be convex. Then

- (1) f is continuous.
- (2) f is differentiable outside a countable subset of (a, b) .
- (3) If $a < x < x' < b$ then $f'_+(x) \leq f'_-(x')$. (So ∂f is monotone.)

Remark 1. (2) does not generalise to the case in which the domain of f is a compact and convex subset of \mathbf{R}^m for $m \geq 2$. It remains true that a convex function is differentiable a.e.

¹These are, respectively, the right- and left-hand derivatives of f at x .

A very useful property of convex functions is that they admit a fundamental theorem of calculus. A proof can be found in Rockafellar (1970, p. 232).

Theorem 1 (convex FToC). Let $f, g \in \mathbf{R}^{[a,b]}$. The following are equivalent:

- (1) f is convex, and $g \in \partial f$.
- (2) g is increasing, and $f(x) = f(a) + \int_a^x g$ for each $x \in [a, b]$.²

Corollary 1. Let $g : [a, b] \rightarrow \mathbf{R}$.

- (1) g is increasing iff $g \in \partial f$ for some convex $f : [a, b] \rightarrow \mathbf{R}$.
- (2) If $g \in \partial f$ for $f : [a, b] \rightarrow \mathbf{R}$ convex, then

$$f(x) - f(x') = \int_{x'}^x g \quad \text{for every } x, x' \in [a, b].$$

Proof. If there is f convex such that $g \in \partial f$, then g is increasing by Theorem 1. If g is increasing, define $f(x) := f(a) + \int_a^x g$; then f is convex and $g \in \partial f$ by Theorem 1. Part (2) is immediate from Theorem 1. ■

Corollary 1 is certainly not as neat as the full FToC. I present it because it is sufficient for our purposes and because unlike the FToC, it generalises directly to the multidimensional case (see section 1.3.1).

1.1.2 Environment

Begin with the case of a single buyer. The buyer has type $t \in \mathcal{T}$, where $\mathcal{T} = [t, \bar{t}] \subseteq \mathbf{R}$. Her type is unknown to the principal, who views it as a random variable T distributed with CDF F . If the agent is of type t and pays transfer p to the principal, she earns utility $t - p$ if she gets the object and $-p$ if she does not. (So she is risk-neutral and has a type-independent outside option.) She maximises expected utility.

A direct mechanism in this setting is a pair of functions $g : \mathcal{T} \rightarrow [0, 1]$ and $p : \mathcal{T} \rightarrow \mathbf{R}$ such that if the agent reports her type as t , she wins the object with probability $g(t)$ and pays transfer $p(t)$. (Note that she pays the transfer $p(t)$ regardless of whether she gets the object.) So an agent of type t who reports her type as t' earns a payoff of $g(t')t - p(t')$. We call g the allocation rule and p the transfer rule.

²That's a Riemann integral. Monotone functions are always Riemann-integrable.

Definition 2. Let (g, p) be a direct mechanism.

(1) (g, p) is incentive-compatible (IC) iff

$$g(t)t - p(t) \geq g(t')t - p(t') \quad \text{for every } t, t' \in \mathcal{T}.$$

(2) (g, p) is individually rational (IR) iff $g(t)t - p(t) \geq 0$ for every $t \in \mathcal{T}$.

(3) g is implementable iff there is a $p : \mathcal{T} \rightarrow \mathbf{R}$ s.t. (g, p) is IC.

As usual, a revelation principle applies, so that we can restrict attention to direct mechanisms that satisfy IC and IR. We'd like to characterise the set of IC and IR direct mechanisms, and then to find the mechanism in that set which maximises expected revenue $\mathbf{E}(p(T))$ for the principal.

1.1.3 Characterisation of IC

In this section, we give two results that sharply characterise the set of IC direct mechanisms. The Spence–Mirrlees theorem shows that an allocation rule g is implementable iff it is increasing. The Mirrlees envelope theorem shows that, up to a constant, any IC mechanism implementing g must have the same payments and payoffs.

The key to proving these results is to translate implementability into a statement about the subdifferential of a convex functions; once that's done, the fundamental theorem of calculus for convex functions (Theorem 1) does the rest.

Lemma 1. $g : \mathcal{T} \rightarrow [0, 1]$ is implementable iff $g \in \partial\phi$ for some convex $\phi : \mathcal{T} \rightarrow \mathbf{R}$.

Proof. Suppose that g is implementable. Then there is $p : \mathcal{T} \rightarrow \mathbf{R}$ such that (g, p) is IC. Define ϕ from g and p as

$$\phi(t) := \max_{t' \in \mathcal{T}} [g(t')t - p(t')] \quad \text{for each } t \in \mathcal{T}.$$

Since ϕ is the pointwise maximum of convex functions, it is convex. Moreover, for every $t, t' \in \mathcal{T}$,

$$\begin{aligned} \phi(t') &\geq g(t)t' - p(t) \\ &= [g(t)t - p(t)] + g(t)(t' - t) \\ &= \phi(t) + g(t)(t' - t), \end{aligned}$$

where the first inequality used the definition of ϕ , and the second used IC. So $g \in \partial\phi$ for ϕ convex.

Now suppose that $g \in \partial\phi$ for some $\phi : \mathcal{T} \rightarrow \mathbf{R}$ convex. Define $p : \mathcal{T} \rightarrow \mathbf{R}$ from g and ϕ by

$$p(t) := g(t)t - \phi(t) \quad \text{for each } t \in \mathcal{T}.$$

Then for every $t, t' \in \mathcal{T}$,

$$\begin{aligned} g(t)t - p(t) &= \phi(t) \\ &\geq \phi(t') + g(t')(t - t') \\ &= [g(t')t' - p(t')] + g(t')(t - t') \\ &= g(t')t - p(t'), \end{aligned}$$

where the first and second equalities used the definition of p , and the inequality used $g(t') \in \partial\phi(t')$. So (g, p) is IC, hence g is implementable. ■

Since the subdifferentials of convex functions are precisely the increasing functions, we obtain the following.

Theorem 2 (Spence (1974), Mirrlees (1976)). g is implementable iff it is increasing.

Proof. Immediate from Lemma 1 and part (1) of Corollary 1 (p. 7). ■

The function ϕ that we constructed in the first part of the proof of Lemma 1 will continue to be useful: for a mechanism (g, p) , let $\phi^{g,p} : \mathcal{T} \rightarrow \mathbf{R}$ be given by

$$\phi^{g,p}(t) := \max_{t' \in \mathcal{T}} [g(t')t - p(t')] \quad \text{for each } t \in \mathcal{T}.$$

$\phi^{g,p}(t)$ is the payoff that type t gets from the mechanism (g, p) . $\phi^{g,p}$ is obviously convex by construction, and we showed in the proof of Lemma 1 that $g \in \partial\phi^{g,p}$ if (g, p) is IC.

Since a convex function is determined up to a constant by its subdifferential, we have the following envelope theorem:

Theorem 3 (Mirrlees (1971)). If (g, p) is IC then

$$\phi^{g,p}(t) - \phi^{g,p}(t') = \int_{t'}^t g \quad \text{for every } t, t' \in \mathcal{T}.$$

Proof. Immediate from Lemma 1 and part (2) of Corollary 1 (p. 7). ■

The Mirrlees envelope theorem delivers the familiar payoff and revenue equivalence results:

Corollary 2 (payoff & revenue equivalence). If (g, p) and (g, p') are IC then

$$\phi^{g,p} - \phi^{g,p'} = p' - p = k$$

for some $k \in \mathbf{R}$.

It's obvious from the Mirrlees envelope theorem that $\phi^{g,p}$ is increasing for any IC (g, p) . Hence

Corollary 3. If (g, p) is IC, then it is IR iff $\phi^{g,p}(\underline{t}) \geq 0$.

It follows that IR imposes no constraints on implementability: if (g, p) is IC but not IR, then $(g, p - \phi^{g,p}(\underline{t}))$ is also IC, and is IR by construction.

Remark 2. It's important for the envelope formula that IC hold on a connected set (in our case an interval $\mathcal{T} = [\underline{t}, \bar{t}]$). When there are gaps, our fundamental theorem of calculus is inapplicable, and results such as revenue equivalence may fail.

1.1.4 Optimal mechanism

By the envelope formula, if (g, p) is IC and IR then type t pays

$$p(t) = g(t)t - \int_{\underline{t}}^t g - \phi^{g,p}(\underline{t}),$$

where $\phi^{g,p}(\underline{t})$ is any nonnegative number. Clearly revenue is maximised by setting $\phi^{g,p}(\underline{t}) = 0$, so that

$$p(t) = g(t)t - \int_{\underline{t}}^t g. \tag{1}$$

Expected revenue as a function of g alone is then

$$R(g) := \mathbf{E} \left(g(T)T - \int_{\underline{t}}^T g \right).$$

The principal's problem is to maximise $R(g)$ subject to $g \in \mathcal{M}_1$, where \mathcal{M}_1 denotes the set of increasing functions $\mathcal{T} \rightarrow [0, 1]$.

R is evidently linear:

$$R(\alpha g + \alpha' g') = \alpha R(g) + \alpha' R(g').$$

\mathcal{M}_1 is clearly convex; it is also compact in the product topology.³ Since we're maximising a linear function on a compact and convex constraint set, the maximum is attained at an extreme point of the constraint set. The extreme points of \mathcal{M}_1 are precisely the functions of the form

$$g = \mathbf{1}_{[t^*, \bar{t}]} \quad \text{or} \quad g = \mathbf{1}_{(t^*, \bar{t}]} \quad \text{for some } t^* \in \mathcal{T}.$$

So we've shown that the optimal mechanism is deterministic.⁴ Further, we've shown that the object is assigned according to a threshold rule. By (1), all types below t^* pay $p(t) = 0$, and types above t^* pay

$$p(t) = t - \int_{t^*}^t dt' = t^*.$$

So the optimal mechanism is a posted price t^* .

Clearly expected revenue from posting the price t^* is

$$[1 - F(t^*)] t^*.$$

(Recall that F is the CDF of the agent's type T .) The optimal price is the one that maximises this expression. When the CDF F has a density f , the optimal price solves the first-order condition

$$1 - F(t^*) - f(t^*) t^* = 0.$$

³ $[0, 1]^{\mathcal{T}}$ is compact by Tychonoff's theorem, and \mathcal{M}_1 is closed in $[0, 1]^{\mathcal{T}}$, hence compact.

⁴Not a trivial result. It doesn't hold with multiple objects. It relies heavily on linearity and the one-dimensional structure.

1.2 Several agents, one-dimensional independent types

1.2.1 Environment

Now suppose that there are n agents indexed by $i \in \{1, \dots, n\}$. Their types are $t = (t_1, \dots, t_n) \in \prod_{i=1}^n \mathcal{T}_i =: \mathcal{T}$, where each $\mathcal{T}_i = [\underline{t}_i, \bar{t}_i]$ is an interval of \mathbf{R} . The principal views types as independently distributed random variables $\{T_i\}$ with marginal CDFs $\{F_i\}$. Preferences are as before.

A direct mechanism is a pair (g, p) of an allocation rule $g : \mathcal{T} \rightarrow [0, 1]^n$ such that $\sum_{i=1}^n g_i(t) \leq 1$ and a transfer rule $p : \mathcal{T} \rightarrow \mathbf{R}^n$.

1.2.2 Bayesian implementation

Consider Bayesian implementation: we design a game and assume that agents play a Bayes–Nash equilibrium. Define, for each i ,

$$G_i(t_i) := \mathbf{E}(g_i(t_i, T_{-i})) \quad \text{and} \quad P_i(t_i) := \mathbf{E}(p_i(t_i, T_{-i})) \quad \text{for each } t_i \in \mathcal{T}_i.$$

Definition 3. Let (g, p) be a direct mechanism.

- (1) (g, p) is Bayes incentive-compatible (BIC) iff for each i ,

$$G_i(t_i)t_i - P_i(t_i) \geq G_i(t'_i)t_i - P_i(t'_i) \quad \forall t_i, t'_i \in \mathcal{T}_i.$$

- (2) (g, p) is interim individually rational (IIR) iff $G_i(t_i)t_i - P_i(t_i) \geq 0$ for every $t_i \in \mathcal{T}_i$, for each agent i .
- (3) g is implementable in Bayes–Nash iff there is a $p : \mathcal{T} \rightarrow \mathbf{R}^n$ such that (g, p) is BIC.

An example of a BIC and IIR direct mechanism is the second-price sealed-bid auction: g assigns the object to the highest bidder (with uniform tie-breaking), and p requires the winner to pay the second-highest bid while the losers pay nothing.

As with a single agent, a revelation principle applies, so that we may restrict attention to direct mechanisms that are BIC and IIR.

When types are independent, we can characterise BIC by (trivially) extending our characterisation results for IC in the one-agent setting. Corollary 4 shows that an allocation rule g is implementable in Bayes–Nash iff it is increasing in an appropriate sense. Corollary 5 shows that, up to a family of constants, any BIC mechanism implementing g must have the same payments and payoffs.

The extension is trivial because of the following observation.

Observation 1. A multi-agent mechanism (g, p) is BIC iff each of the single-agent mechanisms (G_i, P_i) are IC in the sense of section 1.1.2.

Corollary 4. g is implementable in Bayes–Nash iff each G_i is increasing.

Proof. Immediate from the Spence–Mirrlees theorem (p. 9). ■

For a mechanism (g, p) , let $\Phi_i^{g,p} : \mathcal{T}_i \rightarrow \mathbf{R}$ be given by

$$\Phi_i^{g,p}(t_i) := \max_{t'_i \in \mathcal{T}_i} [G_i(t'_i)t_i - P_i(t'_i)] \quad \text{for each } t_i \in \mathcal{T}_i.$$

$\Phi_i^{g,p}(t_i)$ is the (interim) payoff that type t_i of agent i earns in the mechanism, exactly analogous to $\phi^{g,p}$ in section 1.1.3.

Corollary 5. If (g, p) is BIC then for each i ,

$$\Phi_i^{g,p}(t'_i) - \Phi_i^{g,p}(t_i) = \int_{t_i}^{t'_i} G_i \quad \text{for every } t_i, t'_i \in \mathcal{T}_i.$$

Proof. Immediate from the Mirrlees envelope theorem (p. 9). ■

Hence:

Corollary 6 (payoff & revenue equivalence). If (g, p) and (g, p') are BIC then for each i ,

$$\Phi_i^{g,p} - \Phi_i^{g,p'} = P'_i - P_i = k_i$$

for some $k_i \in \mathbf{R}$.

Corollary 7. If (g, p) is BIC, then it is IIR iff $\Phi_i^{g,p}(t_i) \geq 0$ for each i .

When types are interdependent, the results above break down. The natural modifications to the definitions of G_i and P_i in this case are

$$G_i(t_i) := \mathbf{E}(g_i(t_i, T_{-i}) | T_i = t_i) \quad \text{and} \quad P_i(t_i) := \mathbf{E}(p_i(t_i, T_{-i}) | T_i = t_i).$$

for each $t_i \in \mathcal{T}_i$. But then BIC reads

$$G_i(t_i)t_i - P_i(t_i) \geq \mathbf{E}(g_i(t'_i, T_{-i}) | T_i = t_i) t'_i - \mathbf{E}(p_i(t'_i, T_{-i}) | T_i = t_i)$$

for each $t_i, t'_i \in \mathcal{T}_i$. The expectations on the right-hand side are not equal to $G_i(t'_i)$ and $P_i(t'_i)$: when i 's type is t_i but she pretends to be of type t'_i , her information is nevertheless $\{T_i = t_i\}$. BIC is therefore no longer equivalent to IC of the single-agent mechanisms $\{(G_i, P_i)\}$.

1.2.3 The Myerson–Satterthwaite theorem

The envelope formula can be used to derive many economically interesting results, one of which is the topic of this section. Consider the one-object-multiple agents setting again, this time with just two agents. One agent is a buyer: her outside option is not having the object, so her IIR constraint is $G_b(t_b)t_b - P_b(t_b) \geq 0$ as before. The other agent is a seller: her outside option is to keep the object and get no money, so her IIR constraint is instead $G_s(t_s)t_s - P_s(t_s) \geq t_s$. These IIR constraints capture the fact that trade is voluntary. Assume that $\underline{t}_i \geq 0$ for both agents, so that they both value the object. This environment is called bilateral trade.

A mechanism (g, p) is (ex-post) efficient iff $g_b(t) + g_s(t) = 1$ for every $t \in \mathcal{T}$ and $g_i(t_i, t_j) = 1$ whenever $t_i > t_j$. Mechanism (g, p) is (ex-post, weakly) budget-balanced iff $p_b(t) + p_s(t) \geq 0$ for every $t \in \mathcal{T}$.

Theorem 4 (Myerson and Satterthwaite (1983)). For bilateral trade with $\bar{t}_s > \underline{t}_b$, $\mathbf{P}(T_b > T_s) > 0$ and $\mathbf{P}(T_s > T_b) > 0$, there exists no direct mechanism that is efficient, BIC, IIR and budget-balanced.

In English: whenever gains from trade are possible but not certain, efficiency and budget balance are impossible if trade is voluntary. Any budget-balanced trading procedure will lead to misallocation with positive probability.

Proof. Take an arbitrary efficient, BIC and IIR direct mechanism (g, p) ; we will show that it is not budget-balanced. Wlog, let g break ties in favour of the buyer, so that $g_b(t_b, t_s) = 1$ when $t_b = t_s$. Then $G_b(t_b) = \mathbf{P}(t_b \geq T_s)$. We know that a second-price auction is BIC and has interim expected payment $\mathbf{E}(\mathbf{1}(t_b \geq T_s)T_s)$, so by revenue equivalence $P_b(t_b) = \mathbf{E}(\mathbf{1}(t_b \geq T_s)T_s) - k_b$ for some $k_b \in \mathbf{R}$. IIR for \underline{t}_b requires

$$P_b(\underline{t}_b) \leq \mathbf{P}(t_b \geq T_s)\underline{t}_b.$$

It follows that P_b is bounded above (pointwise) by

$$\begin{aligned} \bar{P}_b(t_b) &:= \mathbf{E}(\mathbf{1}(t_b \geq T_s)T_s) - \mathbf{E}(\mathbf{1}(t_b \geq T_s)T_s) + \mathbf{P}(t_b \geq T_s)\underline{t}_b \\ &= \mathbf{E}(\mathbf{1}(t_b \geq T_s > \underline{t}_b)T_s) + \mathbf{P}(t_b \geq T_s)\underline{t}_b \\ &= \mathbf{E}(\mathbf{1}(t_b \geq T_s)[T_s \vee \underline{t}_b]). \end{aligned}$$

This says that revenue from the buyer is maximised by running a modified second-price auction in which the buyer is charged at least her minimum valuation \underline{t}_b when she wins (even if the second bid T_s is $< \underline{t}_b$).

Analogously, $G_s(t_s) = \mathbf{P}(t_s > T_b)$, $P_s(t_s) = \mathbf{E}(\mathbf{1}(t_s > T_b)T_b) - k_s$ by revenue equivalence to a second-price auction, and

$$P_s(\bar{t}_s) \leq \mathbf{P}(\bar{t}_s > T_b)\bar{t}_s - \bar{t}_s = -\mathbf{P}(\bar{t}_s \leq T_b)\bar{t}_s$$

by IIR for \bar{t}_s . Hence P_s is pointwise bounded above by

$$\begin{aligned} \bar{P}_s(t_s) &:= \mathbf{E}(\mathbf{1}(t_s > T_b)T_b) - \mathbf{E}(\mathbf{1}(\bar{t}_s > T_b)T_b) - \mathbf{P}(\bar{t}_s \leq T_b)\bar{t}_s \\ &= -\mathbf{E}(\mathbf{1}(\bar{t}_s > T_b \geq t_s)T_b) - \mathbf{P}(\bar{t}_s \leq T_b)\bar{t}_s \\ &= -\mathbf{E}(\mathbf{1}(T_b \geq t_s)[T_b \wedge \bar{t}_s]). \end{aligned}$$

This payment corresponds to a modified second-price auction in which the seller's remuneration is capped at \bar{t}_s .

It follows that expected total revenue is bounded above by

$$\begin{aligned} \mathbf{E}(\bar{P}_b(T_b) + \bar{P}_s(T_s)) &= \mathbf{E}(\mathbf{1}(T_b \geq T_s)([T_s \vee \underline{t}_b] - [T_b \wedge \bar{t}_s])) \\ &= \mathbf{E}([T_s \vee \underline{t}_b] - [T_b \wedge \bar{t}_s] | T_b \geq T_s) \times \mathbf{P}(T_b \geq T_s), \end{aligned}$$

which is strictly negative since $\mathbf{P}(T_b \geq T_s) > 0$, $\mathbf{P}(T_b > T_s | T_b \geq T_s) \geq \mathbf{P}(T_b > T_s) > 0$, and $\underline{t}_b < \bar{t}_s$. So (g, p) runs a deficit in expectation, hence cannot be budget-balanced. \blacksquare

1.2.4 Strategy-proof implementation

So far, our notion of incentive compatibility for multiple agents has been BIC, which requires that whenever agents $-i$ report their types truthfully, each type t_i of agent i wishes to report truthfully. A more stringent (perhaps more compelling) notion of incentive compatibility requires instead that no matter what agents $-i$ report, each type t_i wishes to report truthfully. A direct mechanism (g, p) with this property is called dominant-strategy incentive-compatible (or strategy-proof). More formally,

Definition 4. Let (g, p) be a direct mechanism.

- (1) (g, p) is dominant-strategies incentive-compatible (DSIC) iff for each i and each $t_{-i} \in \mathcal{T}_{-i}$,

$$g_i(t_i, t_{-i})t_i - p_i(t_i, t_{-i}) \geq g_i(t'_i, t_{-i})t_i - p_i(t'_i, t_{-i}) \quad \forall t_i, t'_i \in \mathcal{T}_i.$$

- (2) g is implementable in dominant strategies iff there is a $p : \mathcal{T} \rightarrow \mathbf{R}^n$ such that (g, p) is DSIC.

Just as for Bayesian implementation, there is a revelation principle for strategy-proof implementation that justifies restricting attention to DSIC and IIR direct mechanisms. It should be obvious that the distribution of types is irrelevant for strategy-proof implementation; in particular, our maintained assumption that types are independent is not needed.

Of course there are many mechanisms which are BIC but not DSIC. But as we'll see in section 1.2.9, for any BIC mechanism (g, p) , there is a DSIC mechanism (g', p') that has the same 'interim allocation' $G'_i = G_i$, and hence the same interim payoffs and expected revenue. So provided you are interested only in the interim allocation (e.g. because you care about interim payoffs or expected revenue), you get DSIC for free! As an example, consider auctions: all BIC auctions are interim allocation-equivalent to the second-price auction, which is DSIC.

DSIC can be characterised in a similar way to BIC by using our characterisation results for single-agent IC. An important difference is that the characterisation of DSIC does not depend on how types are distributed; in particular, types are not required to be independent.

Observation 2. A multi-agent mechanism (g, p) is DSIC iff each of the single-agent mechanisms $(g_i(\cdot, t_{-i}), p_i(\cdot, t_{-i}))$ are IC in the sense of section 1.1.2.

Corollary 8. g is implementable in dominant strategies iff each $g_i(\cdot, t_{-i})$ is increasing.

Proof. Immediate from the Spence–Mirrlees theorem (p. 9). ■

For a mechanism (g, p) , let $\phi_i^{g,p} : \mathcal{T} \rightarrow \mathbf{R}$ be given by

$$\phi_i^{g,p}(t_i, t_{-i}) := \max_{t'_i \in \mathcal{T}_i} [g_i(t'_i, t_{-i})t_i - p_i(t'_i, t_{-i})] \quad \text{for each } (t_i, t_{-i}) \in \mathcal{T}.$$

$\phi_i^{g,p}(t_i, t_{-i})$ is the (ex-post) payoff that type t_i of agent i earns in the mechanism when the others report t_{-i} .

Corollary 9. If (g, p) is DSIC then for each i and $t_{-i} \in \mathcal{T}_{-i}$,

$$\phi_i^{g,p}(t'_i, t_{-i}) - \phi_i^{g,p}(t_i, t_{-i}) = \int_{t_i}^{t'_i} g_i(\cdot, t_{-i}) \quad \text{for every } t_i, t'_i \in \mathcal{T}_i.$$

Proof. Immediate from the Mirrlees envelope theorem (p. 9). ■

Therefore

Corollary 10 (payoff & revenue equivalence). If (g, p) and (g, p') are DSIC then for each i and $t_{-i} \in \mathcal{T}_{-i}$,

$$\phi_i^{g,p}(\cdot, t_{-i}) - \phi_i^{g,p'}(\cdot, t_{-i}) = p'_i(\cdot, t_{-i}) - p_i(\cdot, t_{-i}) = k_{i,t_{-i}}$$

for some $k_{i,t_{-i}} \in \mathbf{R}$.

1.2.5 Equilibrium via revenue equivalence

An easy(ish) way of solving for the Bayes–Nash equilibrium of a difficult auction is to find an easier-to-solve auction that implements the same g , and then to use revenue equivalence. In this section, we’ll solve for the Bayes–Nash equilibrium of a first-price auction via the easier-to-solve second-price auction.

Consider again the single-object-multiple-agents setting. Assume that all agents’ types are drawn from the same distribution, and that its CDF F is differentiable and strictly increasing on $[\underline{t}, \bar{t}]$.

In a second-price auction, we have $g_i(t_i, t_{-i}) = 1$ whenever $t_i > \max_{j \neq i} t_j$ and $p_i(t_i, t_{-i}) = g_i(t_i, t_{-i}) \max_{j \neq i} t_j$. Hence

$$G_i(t_i) = \mathbf{P} \left(t_i \geq \max_{j \neq i} T_j \right) \quad \text{and} \quad P_i(t_i) = \mathbf{E} \left(\mathbf{1} \left(t_i \geq \max_{j \neq i} T_j \right) \max_{j \neq i} T_j \right).$$

Moreover, $G_i(\underline{t}) = P_i(\underline{t}) = 0$.

We look for a symmetric and monotone BNE of the first-price auction, meaning that every player plays an increasing bid function $b : [\underline{t}, \bar{t}] \rightarrow \mathbf{R}_+$. In that case, $t_i > \max_{j \neq i} t_j$ implies that agent i wins the good, so such an equilibrium would implement the same g_i as the second-price auction. Moreover, it implies that an agent of type \underline{t} wins with probability zero, and so pays nothing in expectation. Hence by revenue equivalence, the interim expected payments in such an equilibrium of the first-price auction must be the $\{P_i\}_i$ that we derived for the second-price auction. The interim expected payments in the first-price auction are

$$\mathbf{P} \left(b(t_i) \geq \max_{j \neq i} b(T_j) \right) b(t_i) = \mathbf{P} \left(t_i \geq \max_{j \neq i} T_j \right) b(t_i).$$

Revenue equivalence therefore tells us that

$$\mathbf{P} \left(t_i \geq \max_{j \neq i} T_j \right) b(t_i) = \mathbf{E} \left(\mathbf{1} \left(t_i \geq \max_{j \neq i} T_j \right) \max_{j \neq i} T_j \right) \quad \forall t_i \in [\underline{t}, \bar{t}],$$

or

$$b(t_i) = \mathbf{E} \left(\max_{j \neq i} T_j \mid \max_{j \neq i} T_j < t_i \right) \quad \forall t_i \in [\underline{t}, \bar{t}].$$

We're not quite done yet. Since the direct mechanism (g, p) is BIC, we have established that in the first-price auction, no type t_i wishes to bid $b(t'_i)$ instead of $b(t_i)$ when everyone else is playing b . But this does not rule out deviations to some $b' \notin b([\underline{t}, \bar{t}])$. (The revelation principle is somewhat subtle!) And of course we need to rule out all deviations to conclude that everyone playing b is a Bayes–Nash equilibrium.

In this case, it's easy. b is continuous since F is, and strictly increasing since F is. By the intermediate value theorem, it follows that b has a continuous, strictly increasing inverse b^{-1} . Hence every deviant bid b' in $[b(\underline{t}), b(\bar{t})] = [\underline{t}, b(\bar{t})]$ is the bid of some type, so is unprofitable by BIC. It remains only to show that bids strictly above $b(\bar{t})$ are not profitable deviations, which is obvious since a player can do better (keep the same winning probability while lowering what she pays if she wins) by bidding $b(\bar{t})$ instead.

Finally, a note. While both the second- and first-price auctions studied here implement the same g and are interim revenue-equivalent (have the same P_i s), they do not have the same payment rule p .

1.2.6 Optimal mechanism

We already studied optimal auction design for the case of a single object and a single agent. We found that the optimal mechanism is to post a price t^* : types below t^* pay nothing and don't get the object, whereas types above t^* pay the price and get the object. In this section, we study optimal auction design with multiple bidders, following Myerson (1981). We impose only BIC, though we'll see that the optimal BIC mechanism is in fact DSIC.

Let's begin by attempting to replicate the clean argument that we used for the single-agent case. (Preview: it won't work.) Let \mathcal{M}_n denote the set of functions $g : \mathcal{T} \rightarrow [0, 1]^n$ such that $\sum_{i=1}^n g_i(t) \leq 1$ for every $t \in \mathcal{T}$ and such that G_i is increasing for each i . We've seen that g is implementable in Bayes–Nash exactly if $g \in \mathcal{M}_n$.

Given g , the envelope formula tells us that a BIC mechanism (g, p) must satisfy

$$P_i(t_i) = G_i(t_i)t_i - \int_{\underline{t}_i}^{t_i} G_i - \Phi_i(\underline{t}_i)$$

for every i and $t_i \in \mathcal{T}_i$. Subject to IIR, revenue is clearly maximised by letting $\Phi_i(\underline{t}_i) = 0$. We can therefore write expected revenue as a function of g alone:

$$R(g) := \sum_{i=1}^n \mathbf{E} \left(G_i(T_i)T_i - \int_{\underline{t}_i}^{T_i} G_i \right).$$

As in our previous argument, \mathcal{M}_n is convex and compact,⁵ and R is linear. Hence the optimal mechanism must be an extreme point of \mathcal{M}_n . Extreme points of \mathcal{M}_n take only values in $\{0, 1\}^n$, so it follows that the optimal mechanism is deterministic. But the observation that the optimal mechanism is an extreme point of \mathcal{M}_n isn't as useful as in the single-bidder case because there are a lot of extreme points of \mathcal{M}_n , many of them strange. For example, for $n = 2$, one extreme point of \mathcal{M}_n is the function that always assigns the good to either no-one or to agent 1, with the decision depending (perhaps non-monotonically) on the report of agent 2.

Fortunately, the problem has enough structure that we can proceed in a more specialised manner. Let F_i denote the CDF of the distribution of T_i ,

⁵ $([0, 1]^n)^{\mathcal{T}}$ is compact in the product topology by Tychonoff's theorem, and \mathcal{M}_n is closed in $([0, 1]^n)^{\mathcal{T}}$, hence compact.

and assume that it has a strictly positive density f_i . Rewrite

$$\begin{aligned}
R(g) &= \sum_{i=1}^n \mathbf{E} \left(G_i(T_i)T_i - \int_{\underline{t}_i}^{T_i} G_i(t'_i) dt'_i \right) \\
&= \sum_{i=1}^n \int_{\mathcal{T}_n} \cdots \int_{\mathcal{T}_1} \left[G_i(t_i)t_i - \int_{\underline{t}_i}^{t_i} G_i(t'_i) dt'_i \right] \left(\prod_{j=1}^n f_j(t_j) dt_j \right) \\
&= \sum_{i=1}^n \int_{\underline{t}_i}^{\bar{t}_i} \left[G_i(t_i)t_i - \int_{\underline{t}_i}^{t_i} G_i(t'_i) dt'_i \right] f_i(t_i) dt_i \\
&= \sum_{i=1}^n \left[\int_{\underline{t}_i}^{\bar{t}_i} G_i(t_i)t_i f_i(t_i) dt_i - \int_{\underline{t}_i}^{\bar{t}_i} \left(\int_{\underline{t}_i}^{t_i} G_i(t'_i) dt'_i \right) f_i(t_i) dt_i \right] \\
&= \sum_{i=1}^n \left[\int_{\underline{t}_i}^{\bar{t}_i} G_i(t_i)t_i f_i(t_i) dt_i \right. \\
&\quad \left. - \left[\left(\int_{\underline{t}_i}^{t_i} G_i(t'_i) dt'_i \right) F_i(t_i) \right]_{\underline{t}_i}^{\bar{t}_i} + \int_{\underline{t}_i}^{\bar{t}_i} G_i(t_i) F_i(t_i) dt_i \right] \\
&= \sum_{i=1}^n \left[\int_{\underline{t}_i}^{\bar{t}_i} G_i(t_i)t_i f_i(t_i) dt_i \right. \\
&\quad \left. - \left(\int_{\underline{t}_i}^{\bar{t}_i} G_i(t'_i) dt'_i \right) \cdot 1 + \int_{\underline{t}_i}^{\bar{t}_i} G_i(t_i) F_i(t_i) dt_i \right] \\
&= \sum_{i=1}^n \int_{\underline{t}_i}^{\bar{t}_i} [G_i(t_i)t_i f_i(t_i) - G_i(t_i) + G_i(t_i)F_i(t_i)] dt_i \\
&= \sum_{i=1}^n \int_{\mathcal{T}_i} G_i(t_i) \left[t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} \right] f_i(t_i) dt_i \\
&= \sum_{i=1}^n \int_{\mathcal{T}_n} \cdots \int_{\mathcal{T}_1} g_i(t_1, \dots, t_n) \left[t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} \right] \left(\prod_{j=1}^n f_j(t_j) dt_j \right) \\
&= \mathbf{E} \left(\sum_{i=1}^n g_i(T_1, \dots, T_n) \left[T_i - \frac{1 - F_i(T_i)}{f_i(T_i)} \right] \right),
\end{aligned}$$

where the fifth equality used integration by parts. The function

$$V_i(t_i) := t_i - \frac{1 - F_i(t_i)}{f_i(t_i)}$$

is called the virtual surplus function of agent i . It captures i 's actual valuation

t_i of the object, adjusted by the information rent $(1 - F_i(t_i))/f_i(t_i)$ that she must be paid to report truthfully.

Ignore monotonicity for a moment. Then finding the optimal g is easy: for each profile $t \in \mathcal{T}$, pick the optimal $g(t)$ (‘pointwise maximisation’). That means picking $g(t)$ to maximise $\sum_{i=1}^n g_i(t)V_i(t_i)$, for each $t \in \mathcal{T}$. And that in turn means assigning the object to the agent with the highest virtual surplus $V_i(t_i)$, or to no-one if all agents’ virtual surpluses are negative.

But we can’t ignore monotonicity, or else BIC may be violated. In this section, we will duck that worry by limiting attention to the ‘regular case’ in which each V_i is increasing (i.e. $t_i \mapsto t_i - [1 - F_i(t_i)]/f_i(t_i)$ increasing). This is a restriction on the distributions $\{F_i\}$; a commonly-satisfied sufficient condition is for the hazard rates $\{f_i(\cdot)/[1 - F_i(\cdot)]\}$ be increasing.

We saw that the g that maximises R when we ignore monotonicity is such that i ’s probability of receiving the object is increasing in $V_i(t_i)$. In the regular case, V_i is increasing, and hence i ’s probability of receiving the object is increasing in t_i . So G_i is increasing, hence monotonicity is respected. This is why the regular case is nice: the monotonicity constraint does not bind, and hence we can find the optimal mechanism by maximising pointwise (without worrying about ‘global’ shape restrictions such as monotonicity).

(When regularity fails, monotonicity binds. We then have to find the optimal way of ironing out the nonmonotonocities in the pointwise-optimal mechanism; see section 1.2.8.)

The optimal mechanism for the regular case can evidently be implemented by a second-price auction with reserve prices (one for each agent, determined by her F_i).⁶ So the optimal BIC mechanism is actually DSIC! We get strategy-proofness for free. This generalises: we’ll see in section 1.2.9 that any BIC mechanism has a DSIC counterpart that yields the same ‘interim allocation’ $\{G_i\}$, hence the same payoffs and expected revenue.

Note that the object may not be assigned to anyone, and that when it is assigned it may not be assigned to the agent with the highest reported valuation t_i . (It will be assigned to the agent with the highest reported virtual surplus $t_i - [1 - F_i(t_i)]/f_i(t_i)$.) When types are drawn from the same distribution ($F_i = F$), however, the object is assigned to the highest-valuation type (or to no-one). So in the symmetric case, the optimal mechanism can be implemented by a first-price auction with reserve prices.

⁶Recall that the optimal mechanism for the single-agent case was a posted price. That is clearly a special case of a second-price auction with reserve prices, as we should hope.

1.2.7 The Bulow–Klemperer theorem

We’ve seen that one revenue-maximising way of selling an object to n agents is to run a second-price auction with reserve prices. That such a simple mechanism is optimal is rather pleasing! But we may be unhappy about the fact that the optimal reserve prices depend sensitively on the distributions $\{F_i\}$. In this section we will show that, for the regular case with values drawn from the same distribution F , that a second-price auction *without* reserve prices raises almost as much revenue as an optimal auction, where the ‘almost as’ bound is independent of F . In a word, a second-price auction is approximately optimal, uniformly over F .

We first show that subject to never withholding the object, the second-price auction is optimal. (A result of independent interest.)

Lemma 2. Consider the single-object- n -agents setting with $F_i = F$ for some strictly increasing and differentiable F (with derivative f) such that $t \mapsto t - [1 - F(t)]/f(t)$ is increasing. The second-price auction is revenue-maximising in the class of BIC and IIR mechanisms (g, p) that satisfy $\sum_{i=1}^n g_i(t) = 1 \forall t \in \mathcal{T}$.

Proof. Recall from the derivation of the optimal auction that

$$R(g) = \mathbf{E} \left(\sum_{i=1}^n g_i(T_1, \dots, T_n) V(T_i) \right)$$

where $V(t_i) := t_i - [1 - F(t_i)]/f(t_i)$. Maximising profile-by-profile subject to the constraint $\sum_{i=1}^n g_i(t) = 1$ then yields that you should allocate the good to whoever has the highest $V(t_i)$. Since we’re in the regular case, V is strictly increasing, so the optimal g assigns the object to whoever has the highest value. But this is what the second-price auction implements. Since the second-price auction gives interim expected payoff zero to type t and implements the revenue-optimal g , it is revenue-maximising by revenue equivalence. ■

Write $\Pi^*(n)$ and $\Pi^s(n)$ for the expected revenue from, respectively, an optimal auction and second-price auction (with n bidders).

Theorem 5 (Bulow and Klemperer (1996)). In the single-object- n -agents setting with $F_i = F$ for some strictly increasing and differentiable F (with derivative f) such that $t \mapsto t - [1 - F(t)]/f(t)$ is increasing, $\Pi^*(n) \leq \Pi^s(n+1)$.

The proof is an elegant revealed-preference argument:

Proof. When selling to $n + 1$ agents, one BIC and IIR mechanism is to run an optimal auction for agents $1, \dots, n$, and to give agent $n + 1$ the object for free whenever the optimal auction specifies withholding it. This mechanism obviously raises expected revenue $\Pi^*(n)$. Moreover, it satisfies $\sum_{i=1}^{n+1} g_i(t) = 1$ $\forall t \in \mathcal{T}$, so by Lemma 2 this mechanism raises lower expected revenue than a second-price auction. So $\Pi^*(n) \leq \Pi^s(n + 1)$. ■

The Bulow–Klemperer theorem provides a bound on the gain in expected revenue from running an optimal auction rather than a second-price auction: the bound is the value of adding a bidder. But the value of adding a bidder depends on F , so we haven't yet obtained a bound that is independent of F . To do that, we'll need another lemma that bounds how quickly expected revenue from a second-price auction can grow as the number of bidders is increased independently of F . (Possibly also of independent interest.)

Lemma 3. In the single-object- n -agents setting with $F_i = F$ for some strictly increasing and differentiable F (with derivative f) such that $t \mapsto t - [1 - F(t)]/f(t)$ is increasing, $\Pi^s(n + 1) \leq \frac{n+1}{n-1} \Pi^s(n)$ for every $n \in \mathbf{N}$.

We'll need some concepts to prove this. Let $X^m = (X_1, \dots, X_m)$ be a random m -vector whose components are independently and identically distributed, and let the random m -vector

$$\left(X_{(1)}^m, \dots, X_{(m)}^m \right)$$

be a random permutation of X^m such that $X_{(1)}^m \geq \dots \geq X_{(m)}^m$.^{7,8} $X_{(k)}^m$ is called the k th order statistic of X^m . For a given X^m , define associated random indices $\{I_k\}_{k=1}^m$ by

$$I_k := \max \left\{ i \in \{1, \dots, m\} : X_i = X_{(k)}^m \right\}.$$

(Except in the case of a tie,) I_k is the index of the k th-highest X_i . Facts:

- (1) If the distribution of X_i has CDF F with a density f , then the distribution of $X_{(k)}^m$ has density

$$f_{X_{(k)}^m}(x) = \frac{m!}{(m-k)!(k-1)!} F(x)^{m-k} [1 - F(x)]^{k-1} f(x).$$

⁷I say 'a' random permutation because there may of course be ties. For our purposes, any tie-breaking rule will do since we'll have ties with probability zero and will only care about expectations.

⁸It's actually more common to define $X_{(1)}^m \leq \dots \leq X_{(m)}^m$.

- (2) $\{I_k\}_{k=1}^m$ are distributed independently of the order statistics, and their marginal distributions are uniform on $\{1, \dots, m\}$.

We'll give two proofs of Lemma 3: my short brute-force proof using fact (1), and Eran's elegant proof using fact (2).

Short proof of Lemma 3. Let

$$T^n := (T_1, \dots, T_n) \quad \text{and} \quad T^{n+1} := (T_1, \dots, T_n, T_{n+1}).$$

With n bidders, the winner pays the second-highest valuation, so revenue is $T_{(2)}^n$. If we add an $(n+1)$ th bidder, revenue is $T_{(2)}^{n+1}$. Using the density from fact (1) above, compute

$$\begin{aligned} \Pi^s(n) &= \mathbf{E} \left(T_{(2)}^n \right) \\ &= \int_{\mathbf{R}} n(n-1)F(\tau)^{n-2} [1 - F(x)] f(x) \cdot \tau d\tau \\ &= \frac{n-1}{n+1} \int_{\mathbf{R}} (n+1)nF(\tau)^{n-1} [1 - F(x)] f(x) \cdot F(\tau)^{-1} \cdot \tau d\tau \\ &\geq \frac{n-1}{n+1} \int_{\mathbf{R}} (n+1)nF(\tau)^{n-1} [1 - F(x)] f(x) \cdot \tau d\tau \\ &= \frac{n-1}{n+1} \mathbf{E} \left(T_{(2)}^{n+1} \right) \\ &= \frac{n-1}{n+1} \Pi^s(n+1). \quad \blacksquare \end{aligned}$$

Elegant proof of Lemma 3. Let

$$T^n := (T_1, \dots, T_n) \quad \text{and} \quad T^{n+1} := (T_1, \dots, T_n, T_{n+1}).$$

With n bidders, the winner pays the second-highest valuation, so revenue is $T_{(2)}^n$. If we add an $(n+1)$ th bidder, revenue is $T_{(2)}^{n+1}$. Observe that revenue changes iff T_{n+1} is the highest or second-highest value:

$$T_{(2)}^n = \begin{cases} T_{(3)}^{n+1} & \text{if } T_{n+1} \text{ is the first- or second-highest value} \\ T_{(2)}^{n+1} & \text{otherwise.} \end{cases}$$

Equivalently, letting $\{I_k\}_{k=1}^m$ be the random indices associated with T^{n+1} ,

$$T_{(2)}^n = \begin{cases} T_{(3)}^{n+1} & \text{if } I_1 = n+1 \text{ or } I_2 = n+1 \\ T_{(2)}^{n+1} & \text{otherwise.} \end{cases}$$

Write $E := \{I_1 = n + 1\} \cup \{I_2 = n + 1\}$. By fact (2) above, the order statistics are independent of E , and $\mathbf{P}(E) = 2/(n + 1)$. Hence

$$\begin{aligned}
\Pi^s(n) &= \mathbf{E} \left(T_{(2)}^n \right) \\
&= \mathbf{P}(E^c) \mathbf{E} \left(T_{(2)}^{n+1} \middle| E^c \right) + \mathbf{P}(E) \mathbf{E} \left(T_{(3)}^{n+1} \middle| E \right) \\
&= \frac{n-1}{n+1} \mathbf{E} \left(T_{(2)}^{n+1} \right) + \frac{2}{n+1} \mathbf{E} \left(T_{(3)}^{n+1} \right) \\
&\geq \frac{n-1}{n+1} \mathbf{E} \left(T_{(2)}^{n+1} \right) \\
&= \frac{n-1}{n+1} \Pi^s(n+1). \quad \blacksquare
\end{aligned}$$

Corollary 11. In the single-object- n -agents setting with $F_i = F$ for some strictly increasing and differentiable F (with derivative f) such that $t \mapsto t - [1 - F(t)]/f(t)$ is increasing, $\Pi^*(n) \leq \frac{n+1}{n-1} \Pi^s(n)$.

Proof. $\Pi^*(n) \leq \Pi^s(n+1) \leq \frac{n+1}{n-1} \Pi^s(n)$ by the Bulow–Klemperer theorem (first inequality) and Lemma 3 (second inequality). \blacksquare

This is an F -independent bound on the gains from adding optimal reserve prices to a second-price auction, so we're done.

1.2.8 Ironing

In this section, it will be convenient to adopt a change of variables. Without loss of generality, we may take types to be uniformly distributed on $[0, 1]$, on the understanding that valuations are given by some increasing function of type. In particular, the type τ_i of an agent with valuation t_i is $\tau_i = F_i(t_i)$, so that $d\tau_i = f_i(t_i) dt_i$ and $F_i(T_i) \sim \mathcal{U}[0, 1]$. The correct inverse of this map for our purposes is

$$F_i^{-1}(\tau_i) := \inf \{t'_i \in \mathcal{T}_i : F_i(t'_i) \geq \tau_i\}.$$

(Draw a picture to convince yourself.) The map from types in $[0, 1]$ to virtual surplus is then given by $V_i \circ F_i^{-1}$. Note that since F_i^{-1} is increasing, $V_i \circ F_i^{-1}$ has the same monotonicity properties as V_i . For the remainder of this section, we will forget about valuations, so we'll write $t_i \in [0, 1]$ and $T_i \sim \mathcal{U}[0, 1]$ instead of τ_i or $F_i^{-1}(t_i)$ for types in $[0, 1]$, and $V_i : [0, 1] \rightarrow \mathbf{R}$ instead of

$V_i \circ F_i^{-1}$ for virtual surplus on $[0, 1]$. That means, note, that V_i is now defined

$$\begin{aligned} V_i(t_i) &:= F_i^{-1}(t_i) - \frac{1 - F_i(F_i^{-1}(t_i))}{f_i(F_i^{-1}(t_i))} \\ &= F_i^{-1}(t_i) - \frac{1 - t_i}{f_i(F_i^{-1}(t_i))} \quad \text{for } t_i \in [0, 1]. \end{aligned}$$

On to business. In section 1.2.6, we characterised the optimal way to sell an object to n buyers in the regular case. We proceeded as follows: using the envelope formula and the fact that it is revenue-maximising to let IIR hold with equality for the lowest type, we wrote down expected revenue as a function R of the allocation g alone. An entire page of algebra showed us that

$$R(g) = \mathbf{E} \left(\sum_{i=1}^n G_i(T_i) V_i(T_i) \right) = \mathbf{E} \left(\sum_{i=1}^n g_i(T_1, \dots, T_n) V_i(T_i) \right).$$

We seek the allocation $g : [0, 1]^n \rightarrow [0, 1]^n$ that maximises revenue R subject to feasibility ($\sum_{i=1}^n g_i(t) \leq 1 \forall t \in \mathcal{T}$) and Bayesian implementability ($G_i(\cdot) = \int_{[0, 1]^{n-1}} g_i(\cdot, t_{-i}) dt_{-i}$ increasing $\forall i$). In the ‘regular’ case in which each V_i is increasing, the relaxed program that drops monotonicity is solved by a g that satisfies monotonicity, and we’re done. We will now tackle the ‘irregular’ case in which monotonicity is a binding constraint. The strategy is called ‘ironing’, and originates (as so many things do) with Myerson (1981).⁹

Here’s the basic recipe. For each i , we wish that the function V_i were increasing. Equivalently, we wish that its integral $W_i(t_i) := \int_0^{t_i} V_i$ were convex. There’s a standard way of making a nonconvex function convex: the convexification $\text{vex } f$ of a function $f : [a, b] \rightarrow \mathbf{R}$ is the (pointwise) largest convex function smaller than f . Equivalently,

$$\text{vex } f(x) = \min \{ \lambda f(x') + (1 - \lambda) f(x'') : \lambda x' + (1 - \lambda) x'' = x, \lambda \in [0, 1] \}.^{10}$$

So let’s use $\text{vex } W_i$ in place of W_i , and accordingly define \bar{V}_i to be any selection from the subdifferential $\partial \text{vex } W_i$.¹¹ Since $\text{vex } W_i$ is convex, $\partial \text{vex } W_i$

⁹There is a precedent in Mussa and Rosen (1978).

¹⁰For f defined on a subset of \mathbf{R}^n , $\text{vex } f$ is still defined as the largest convex function smaller than f , but the characterisation of $\text{vex } f$ as a convex combination of f ’s values requires $n + 1$ points. (Think about it. Remember Carathéodory’s theorem.)

¹¹That is, \bar{V}_i is any map $\mathcal{T}_i \rightarrow \mathbf{R}$ that satisfies $\bar{V}_i(t_i) \in \partial \text{vex } W_i(t_i)$ at each $t_i \in \mathcal{T}_i$. Note in particular that $\bar{V}_i = (\text{vex } W_i)'$ wherever $\text{vex } W_i$ is differentiable (a.e. by convexity).

is monotone, so \bar{V}_i is increasing. Myerson (1981) called \bar{V}_i the ‘ironed’ version of V_i , the idea (presumably) being that we’ve ‘ironed out’ all the nonmonotonicities in V_i .

As in section 1.2.6, write \mathcal{M}_n for the set of functions $g : [0, 1]^n \rightarrow [0, 1]^n$ such that $\sum_{i=1}^n g_i(t) \leq 1$ for every $t \in [0, 1]$ and such that G_i is increasing for each i ; this is the feasible set. Write \mathcal{M}'_n for the larger set of functions $g : [0, 1]^n \rightarrow [0, 1]^n$ that satisfy feasibility $\sum_{i=1}^n g_i(t) \leq 1$, but not necessarily monotonicity. Our problem is $\max_{g \in \mathcal{M}_n} R(g)$. The monotonicity constraint embedded in \mathcal{M}_n makes this problem hard, so we’d like to find an alternative problem whose constraint set is \mathcal{M}'_n , but whose solution is the same as that of $\max_{g \in \mathcal{M}_n} R(g)$.

For the regular case, it was very easy: $\max_{g \in \mathcal{M}'_n} R(g)$ is a suitable alternative problem (we didn’t have to change the objective function). For whenever $\{V_i\}$ are all increasing, the maximiser $g^*(R, \mathcal{M}'_n)$ of R on \mathcal{M}'_n lies in \mathcal{M}_n . That is no longer true when some V_i has nonmonotonicities, so a suitable alternative program must have an objective function different from R . Consider

$$\bar{R}(g) := \mathbf{E} \left(\sum_{i=1}^n G_i(T_i) \bar{V}_i(T_i) \right) = \mathbf{E} \left(\sum_{i=1}^n g_i(T_1, \dots, T_n) \bar{V}_i(T_i) \right),$$

the expected revenue in a fictional regular environment with virtual surpluses $\{\bar{V}_i\}$. Since we’re back in the regular case, we know that $g^*(\bar{R}, \mathcal{M}'_n) \in \mathcal{M}_n$. To establish that this is a suitable alternative problem, it remains only to show its solution coincides with the solution to the original problem: $g^*(\bar{R}, \mathcal{M}'_n) = g^*(R, \mathcal{M}_n)$.

We’ll do this in two steps. We first show that the value of the alternative problem is an upper bound on the value of the original problem:

$$R(g) \leq \bar{R} \left(g^* \left(\bar{R}, \mathcal{M}'_n \right) \right) \quad \text{for any } g \in \mathcal{M}_n. \quad (2)$$

We then show that this upper bound is attained at $g^*(\bar{R}, \mathcal{M}'_n)$:

$$R \left(g^* \left(\bar{R}, \mathcal{M}'_n \right) \right) = \bar{R} \left(g^* \left(\bar{R}, \mathcal{M}'_n \right) \right). \quad (3)$$

As in section 1.1.4, let \mathcal{M}_1 denote the set of increasing functions $[0, 1] \rightarrow [0, 1]$. As we noted there, \mathcal{M}_1 is convex and closed, and its extreme points are

$$\text{ext}\mathcal{M}_1 = \left\{ \mathbf{1}_{[t_i, 1]} : t_i \in [0, 1] \right\} \cup \left\{ \mathbf{1}_{(t_i, 1]} : t_i \in [0, 1] \right\}.$$

Observe that for any linear functional $H : \mathcal{M}_1 \rightarrow \mathbf{R}$, $H \geq 0$ holds on \mathcal{M}_1 iff it holds on $\text{ext}\mathcal{M}_1$.

To prove (2), consider the linear functional $H(G_i) := \int_0^1 G_i(\bar{V}_i - V_i)$. $H \geq 0$ on $\text{ext}\mathcal{M}_1$ since for $G_i = \mathbf{1}_{[t_i, 1]}$ or $G_i = \mathbf{1}_{(t_i, 1]}$,

$$\begin{aligned} \int_0^1 G_i \bar{V}_i &= \int_{t_i}^1 \bar{V}_i = \text{vex } W_i(1) - \text{vex } W_i(t_i) \\ &\geq W_i(1) - W_i(t_i) = \int_{t_i}^1 V_i = \int_0^1 G_i V_i, \end{aligned}$$

where the inequality used the fact that $\text{vex } W_i \leq W_i$ everywhere, with equality at the endpoints 0 and 1. Hence $H \geq 0$ holds on \mathcal{M}_1 , i.e.

$$\int_0^1 G_i \bar{V}_i \geq \int_0^1 G_i V_i \quad \text{for any } G_i \in \mathcal{M}_1.$$

It follows that

$$\begin{aligned} \bar{R}(g) &= \sum_{i=1}^n \mathbf{E} \left(G_i(T_i) \bar{V}_i(T_i) \right) = \sum_{i=1}^n \int_0^1 G_i \bar{V}_i \\ &\geq \sum_{i=1}^n \int_0^1 G_i V_i = \sum_{i=1}^n \mathbf{E} \left(G_i(T_i) V_i(T_i) \right) = R(g), \end{aligned}$$

whence we obtain (2):

$$R(g) \leq \bar{R}(g) \leq \max_{g \in \mathcal{M}_n} \bar{R}(g) = \bar{R} \left(g^* \left(\bar{R}, \mathcal{M}'_n \right) \right) \quad \text{for any } g \in \mathcal{M}_n.$$

To derive (3), fix an i and partition $[0, 1]$ as

$$\left\{ [b_i^0, a_i^1], (a_i^1, b_i^1), [b_i^1, a_i^2], \dots, [b_i^{K-1}, a_i^K], (a_i^K, b_i^K), [b_i^K, a_i^{K+1}] \right\}$$

where $b_i^0 := 0$ and $a_i^{K+1} := 1$, with $\{a_i^k, b_i^k\}_{k=1}^K$ chosen so that

$$\begin{aligned} \text{vex } W_i &< W_i && \text{on } (a_i^k, b_i^k) \text{ for each } k \in [1, K+1] \cap \mathbf{N} \\ \text{vex } W_i &= W_i && \text{on } [b_i^k, a_i^{k+1}] \text{ for each } k \in [0, K] \cap \mathbf{N}. \end{aligned}$$

(I wrote the partition as finite to avoid clutter. It need not be finite, but that does not affect the argument.)

Since $\text{vex } W_i = W_i$ on each $[b_i^k, a_i^{k+1}]$, $\bar{V}_i = V_i$ on each $[b_i^k, a_i^{k+1}]$. We can therefore write our linear functional H as

$$H(G_i) = \int_0^1 G_i (\bar{V}_i - V_i) = \sum_{k=1}^K \int_{a_i^k}^{b_i^k} G_i (\bar{V}_i - V_i).$$

Now consider the $\{G_i\}$ that arise from $g^*(\bar{R}, \mathcal{M}'_1)$:

$$G_i^*(t_i) := \mathbf{E} \left(\left[g^* \left(\bar{R}, \mathcal{M}'_1 \right) \right]_i (t_i, T_{-i}) \right) \quad \text{for } t_i \in [0, 1].$$

On every interval (a_i^k, b_i^k) , $\text{vex } W_i < W_i$, hence $\text{vex } W_i$ is linear, hence \bar{V}_i is constant. Since $g^*(\bar{R}, \mathcal{M}'_1)$ allocates to whoever has the highest $\bar{V}_i(t_i)$ and since types are independent, it follows that G_i^* is constant on every interval (a_i^k, b_i^k) . Hence

$$\begin{aligned} H(G_i^*) &= \sum_{k=1}^K G_i^*(a_i^k) \int_{a_i^k}^{b_i^k} (\bar{V}_i - V_i) = \sum_{k=1}^K G_i^*(a_i^k) [\text{vex } W_i - W_i]_{a_i^k}^{b_i^k} \\ &= \sum_{k=1}^K G_i^*(a_i^k) (0 - 0) = 0 \end{aligned}$$

since $\text{vex } W_i = W_i$ at each a_i^k and each b_i^k by construction. So we've established (3).

To sum up, we've obtained a characterisation of the revenue-maximising mechanism that applies to both the regular and irregular cases. The optimal allocation rule $g^*(R, \mathcal{M}_n)$ is as follows: from the virtual values $\{V_i\}$ construct the ironed virtual values $\{\bar{V}_i\}$, and allocate the object to whoever has the highest \bar{V}_i , or to no-one if everyone's ironed virtual values are negative. The corresponding transfers p are given by the envelope formula and the boundary condition that the lowest type's IIR constraint holds with equality. As in the regular case, the optimal mechanism is actually DSIC, and can be implemented by a second-price auction with a reserve price.

Ironing techniques show up in places other than mechanism design. In the monopoly problem, if marginal revenue is not monotonically decreasing then it needs to be ironed. More generally, 'ironing' often shows up in the solutions to optimal control problems. When we study persuasion and repeated games with one-sided incomplete information in section 3.1 below, we'll make heavy use of a 'concavification' operator that is entirely analogous to the convexification operator here.

1.2.9 Equivalence of Bayesian and strategy-proof implementation

We've seen that when we require BIC, we often get DSIC for free. This is not a fluke: Manelli and Vincent (2010) show that in the single-object model with one-dimensional independent types that we've been studying, for any BIC (g, p) , there is a DSIC (g', p') that yields the same interim allocations.

Gershkov, Goeree, Kushnir, Moldovanu and Shi (2013) extend this result to a slightly larger class of problems using a substantially neater argument; we'll follow their approach. The argument relies on (a more general version of) the following result.

Theorem 6 (Gutmann, Kemperman, Reeds and Shepp (1991)). Let X_1 and X_2 be random elements of $[0, 1]$, and let $q : [0, 1]^2 \rightarrow [0, 1]$ satisfy

$$\mathbf{E}(q(x_1, X_2)) = x_1 \quad \text{and} \quad \mathbf{E}(q(X_1, x_2)) = x_2 \quad \forall (x_1, x_2) \in [0, 1]^2. \quad (4)$$

Then there is an increasing $q' : [0, 1]^2 \rightarrow [0, 1]$ which satisfies (4).

Sketch proof. Consider the program

$$\begin{aligned} \min_{q' \in [0, 1]^{[0, 1]^2}} \quad & \mathbf{E} \left(q'(X_1, X_2)^2 \right) \\ \text{s.t.} \quad & \mathbf{E}(q'(x_1, X_2)) = x_1 \quad \forall x_1 \in [0, 1] \\ & \mathbf{E}(q'(X_1, x_2)) = x_2 \quad \forall x_2 \in [0, 1]. \end{aligned}$$

(Ignore measurability and integrability.) The constraint set is nonempty since q satisfies the constraints. Moreover, the constraint set is compact and the objective continuous in the product topology; hence a minimum exists.

It remains only to show that the minimum q' is increasing. Suppose for a contradiction that it is not: $q'(x_1, x_2) > q'(x'_1, x_2)$ for $x_1 < x'_1$. The constraints then imply that there is $x'_2 > x_2$ such that $q'(x_1, x'_2) < q'(x'_1, x'_2)$. Now consider q'' everywhere equal to q' except that

$$\begin{aligned} q''(x_1, x_2) &= q'(x_1, x_2) - \varepsilon, \\ q''(x'_1, x_2) &= q'(x'_1, x_2) + \varepsilon, \\ q''(x_1, x'_2) &= q'(x_1, x'_2) + \varepsilon, \\ \text{and} \quad q''(x'_1, x'_2) &= q'(x'_1, x'_2) - \varepsilon. \end{aligned}$$

q'' satisfies the constraints since q' does. Assume for simplicity that (X_1, X_2) is uniformly distributed on a discrete support; then

$$\begin{aligned} & \mathbf{E} \left(q''(X_1, X_2)^2 \right) - \mathbf{E} \left(q'(X_1, X_2)^2 \right) \\ &= \left[-2q'(x_1, x_2)\varepsilon + \varepsilon^2 \right] + \left[2q'(x'_1, x_2)\varepsilon + \varepsilon^2 \right] \\ &\quad + \left[2q'(x_1, x'_2)\varepsilon + \varepsilon^2 \right] + \left[-2q'(x_1, x_2)\varepsilon + \varepsilon^2 \right] \\ &= -2\varepsilon \left[q'(x_1, x_2) + q'(x'_1, x'_2) - q'(x'_1, x_2) - q'(x_1, x'_2) \right] + 4\varepsilon^2, \end{aligned}$$

which is strictly negative for $\varepsilon > 0$ small. So q'' is feasible and does better than q' ; hence q' cannot have been a minimiser. \blacksquare

Now consider the multiple-agents-single-object environment. We have the following nice result.

Theorem 7 (Manelli and Vincent (2010)). For any BIC mechanism (g, p) , there is a DSIC mechanism (g', p') such that $G_i = G'_i$ for each i .

Proof. For convenience, assume that there are two agents and that the object must be allocated with probability 1. Relaxing the former just requires more notation, and relaxing the latter is straightforward.

Fix a BIC mechanism (g, p) , and write $G_i(\cdot) := \mathbf{E}(g_i(\cdot, T_{-i}))$ for g 's marginals. Define $q : [0, 1]^2 \rightarrow [0, 1]$ by

$$q(r_1, r_2) := g_1 \left(G_1^{-1}(r_1), (1 - G_2)^{-1}(r_2) \right) \quad \forall (r_1, r_2) \in [0, 1]^2,$$

where the inverses are

$$\begin{aligned} G_1^{-1}(r_1) &:= \inf \{t_1 \in \mathcal{T}_1 : G_1(t_1) \geq r_1\} \quad \forall r_1 \in [0, 1] \\ (1 - G_2)^{-1}(r_2) &:= \inf \{t_2 \in \mathcal{T}_2 : 1 - G_2(t_2) \geq r_2\} \quad \forall r_2 \in [0, 1]. \end{aligned}$$

Further define the random variables $R_1 := G_1(T_1)$ and $R_2 := 1 - G_2(T_2)$.

By definition of G_1 and G_2 , we have for each $(r_1, r_2) \in [0, 1]^2$ that

$$\mathbf{E}(q(r_1, R_2)) = \mathbf{E} \left(g_1 \left(G_1^{-1}(r_1), T_2 \right) \right) = G_1 \left(G_1^{-1}(r_1) \right) = r_1$$

and

$$\begin{aligned} \mathbf{E}(q(R_1, r_2)) &= \mathbf{E} \left(g_1 \left(T_1, (1 - G_2)^{-1}(r_2) \right) \right) \\ &= 1 - G_2 \left((1 - G_2)^{-1}(r_2) \right) = r_2. \end{aligned}$$

So q satisfies (4). Hence by Theorem 6, there is q' increasing which also satisfies (4). Define $g'_1 : \mathcal{T}_1 \times \mathcal{T}_2 \rightarrow [0, 1]$ by

$$g'_1(t_1, t_2) := q' \left(G_1(t_1), 1 - G_2(t_2) \right) \quad \forall (t_1, t_2) \in \mathcal{T}_1 \times \mathcal{T}_2,$$

and let $g'_2 := 1 - g'_1$ and $g := (g'_1, g'_2)$. g' thus defined is an allocation rule. Since q' satisfies (4), the marginals of g' are

$$\begin{aligned} G'_1(t_1) &= \mathbf{E}(g'_1(t_1, T_2)) = \mathbf{E}(q'(G_1(t_1), R_2)) = G_1(t_1) \\ \text{and } G'_2(t_2) &= \mathbf{E}(1 - g'_1(T_1, t_2)) = \mathbf{E}(1 - q'(R_1, 1 - G_2(t_2))) = G_2(t_2), \end{aligned}$$

so g' yields the same 'interim allocation' as g . Since q' is increasing in both its arguments, $g'_1(\cdot, t_2)$ and $g'_2(t_1, \cdot)$ are increasing for every $(t_1, t_2) \in \mathcal{T}_1 \times \mathcal{T}_2$, so g' is implementable in dominant strategies.

Let p' be the transfers pinned down by g' via the envelope formula together whatever boundary condition p satisfies; then (g', p') is a DSIC mechanism that achieves the same interim allocation as (g, p) . ■

Corollary 12. For any BIC mechanism (g, p) , there is a DSIC mechanism (g', p') with the same interim payoffs and expected revenue.¹²

¹²Gershkov et al. (2013) prove that this ‘payoff/revenue equivalence’ extends to a somewhat more general setting, and show by example that the ‘interim allocation equivalence’ in the Manelli–Vincent theorem does not.

1.3 Single agent, multidimensional types

1.3.1 Convex analysis on multidimensional domains

The next topic is multidimensional type spaces. This will require us to extend our results from section 1.1.1 about convex functions on one-dimensional domains to multiple dimensions. Proofs, details and related results can be found in Rockafellar (1970), mainly in sections 23 and 24. Though I (and Rockafellar (1970)) state these results for convex domains $\subseteq \mathbf{R}^m$, they all hold in more general convex spaces.¹³

Definition 5. For X convex, $f : X \rightarrow \mathbf{R}$ is convex iff for any $x, x' \in X$ and $\lambda \in [0, 1]$,

$$f(\lambda x + (1 - \lambda)x') \leq \lambda f(x) + (1 - \lambda)f(x').$$

This definition is ‘one-dimensional’ in that it considers, separately, each one-dimensional linear subspace (i.e. the straight line through any two points).

Convexity is still preserved under pointwise maximisation:

Proposition 4. Let X be convex, and let $\{f_\alpha\}$ be a family of convex functions $X \rightarrow \mathbf{R}$. Define $f : X \rightarrow \mathbf{R}$ by $f(x) := \sup_\alpha f_\alpha(x) \forall x \in X$. Then f is convex.

Let $X \subseteq \mathbf{R}^m$ be compact and convex. For $f : X \rightarrow \mathbf{R}$ convex, $x^* \in \mathbf{R}^m$ is a subgradient of f at $x \in X$ iff

$$f(x') \geq f(x) + x^* \cdot (x' - x) \quad \forall x' \in X.$$

This means precisely that the graph of the affine map $x' \mapsto f(x) + x^* \cdot (x' - x)$ is a supporting hyperplane to $\text{epi } f$ at $(x, f(x))$.¹⁴ The subdifferential of a convex function $f : X \rightarrow \mathbf{R}$ is the correspondence $\partial f : X \rightrightarrows \mathbf{R}^m$ such that $\partial f(x)$ is the set of all subgradients of f at x . For $f : X \rightarrow \mathbf{R}$ and $g : X \rightarrow \mathbf{R}^m$, we write ‘ $g \in \partial f$ ’ as shorthand for ‘ $g(x) \in \partial f(x)$ for every $x \in X$ ’.

Proposition 5. Let $X \subseteq \mathbf{R}^m$ be compact and convex, and let $f : X \rightarrow \mathbf{R}$ be convex. Then

- (1) $\partial f(x)$ is nonempty, compact and convex for every $x \in X$.

¹³They certainly all hold in convex Hilbert spaces. Many of them require less structure than that; see e.g. Rockafellar (1966).

¹⁴ $\text{epi } f$ denotes the epigraph of $f : X \rightarrow \mathbf{R}$, defined $\text{epi } f := \{(x, y) \in X \times \mathbf{R} : y \geq f(x)\}$. Note that f is convex iff $\text{epi } f$ is a convex set.

(2) $\nabla_v f(x)$ exists at every $(x, v) \in \text{int } X \times \mathbf{R}^m$.¹⁵

(3) $\nabla_v f(x) = \max_{x^* \in \partial f(x)} x^* \cdot v$ at every $(x, v) \in \text{int } X \times \mathbf{R}^m$.

This result is the multidimensional analogue of Proposition 2 (p. 6). (1) and (2) are basically unchanged, but property (3) is a little different.

Proposition 6. Let $X \subseteq \mathbf{R}^m$ be compact and convex, and let $f : X \rightarrow \mathbf{R}$ be convex. Then

(1) f is continuous.

(2) f is differentiable outside a set of Lebesgue measure zero.

(1) is a direct extension of property (1) in Proposition 3 (p. 6). (2) is similar to, but weaker than, property (2) in Proposition 3.¹⁶

The final property of convex function that we used was the fundamental theorem of calculus (Theorem 1, p. 7). In the present setting, we can't quite get a fundamental theorem of calculus, but we can get close: in particular, a direct extension of Corollary 1 (p. 7) is available. That extension requires a suitable extension of the notion of monotonicity to the present setting.

Definition 6. Let $X \subseteq \mathbf{R}^m$ be compact and convex. $g : X \rightarrow \mathbf{R}^m$ is cyclically monotone iff

$$g(x_0) \cdot (x_1 - x_0) + g(x_1) \cdot (x_2 - x_1) + \cdots + g(x_{n-1}) \cdot (x_n - x_{n-1}) + g(x_n) \cdot (x_0 - x_n) \leq 0$$

for every finite cycle $x_0, x_1, \dots, x_{n-1}, x_n, x_0$.

Theorem 8 (Rockafellar (1966)). Let $X \subseteq \mathbf{R}^m$ be compact and convex, and let $g : X \rightarrow \mathbf{R}^m$.

(1) g is cyclically monotone iff $g \in \partial f$ for some convex $f : X \rightarrow \mathbf{R}$.

(2) If $g \in \partial f$ for $f : X \rightarrow \mathbf{R}$ convex, then

$$f(x') - f(x) = \int_0^1 g(x + \lambda(x' - x)) \cdot (x' - x) d\lambda \quad \text{for every } x, x' \in X.$$

¹⁵ $\nabla_v f(x)$ denotes the directional derivative of f at x in direction v , i.e.

$$\nabla_v f(x) := \lim_{\delta \downarrow 0} \frac{f(x + \delta v) - f(x)}{\delta}.$$

¹⁶For $X \subseteq \mathbf{R}^m$ with $m \geq 2$, there are convex functions $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ with uncountably many points of nondifferentiability.

Direct proofs of both parts can be found in Rockafellar (1970) (pp. 238 and 239–40, respectively). (2) is straightforward:

Proof of part (2). Fix $x, x' \in X$. Define $\phi : [0, 1] \rightarrow \mathbf{R}$ by

$$\phi(\lambda) := f(x + \lambda(x' - x)) \quad \forall \lambda \in [0, 1],$$

and define $\gamma : [0, 1] \rightarrow \mathbf{R}$ by

$$\gamma(\lambda) := g(x + \lambda(x' - x)) \cdot (x' - x) \quad \forall \lambda \in [0, 1].$$

ϕ is convex since f is. For all $\lambda, \lambda' \in [0, 1]$, we have

$$\begin{aligned} \phi(\lambda') &= f(x + \lambda'(x' - x)) \\ &\geq f(x + \lambda(x' - x)) \\ &\quad + g(x + \lambda(x' - x)) \cdot ([x + \lambda'(x' - x)] - [x + \lambda(x' - x)]) \\ &= f(x + \lambda(x' - x)) + g(x + \lambda(x' - x)) \cdot (x' - x)(\lambda' - \lambda) \\ &= \phi(\lambda) + \gamma(\lambda)(\lambda' - \lambda), \end{aligned}$$

where the inequality used $g(x + \lambda(x' - x)) \in \partial f(x + \lambda(x' - x))$; hence $\gamma \in \partial \phi$. It follows by the fundamental theorem of calculus for convex functions on one-dimensional domains (Theorem 1, p. 7) that

$$f(x') - f(x) = \phi(1) - \phi(0) = \int_0^1 \gamma = \int_0^1 g(x + \lambda(x' - x)) \cdot (x' - x) d\lambda. \quad \blacksquare$$

We will prove part (1) via a result from linear programming, following Vohra (2011). To that end, consider a weighted, directed graph (X, w) , where X are the vertices and $w : X^2 \rightarrow \mathbf{R}$ the edge weights. $f : X \rightarrow \mathbf{R}$ is a potential function for (X, w) iff $f(x') \geq f(x) + w(x, x')$ for every $x, x' \in X$. This is a system of linear inequalities with a very special structure; as a result, we have the following clean characterisation of when they have a solution.

Theorem 9. A graph (X, w) admits a potential function iff

$$w(x_0, x_1) + w(x_1, x_2) + \cdots + w(x_{n-1}, x_n) + w(x_n, x_0) \leq 0 \quad (5)$$

for every finite cycle $x_0, x_1, \dots, x_{n-1}, x_n, x_0$.

Proof of part (1) of Theorem 8. Fix $g : X \rightarrow \mathbf{R}^m$, and define $w : X^2 \rightarrow \mathbf{R}$ by $w(x, x') := g(x) \cdot (x' - x)$ for each $x, x' \in X$. Obviously w satisfies (5) iff g is cyclically monotone.

Suppose that g is cyclically monotone, so that w satisfies (5). Then (X, w) admits a potential $f : X \rightarrow \mathbf{R}$ by Theorem 9. By definition of a potential,

$$f(x') \geq f(x) + w(x, x') = f(x) + g(x) \cdot (x' - x) \quad \forall x, x' \in X;$$

hence $g \in \partial f$. Moreover, f must be convex since ∂f is nonempty-valued.

Now suppose that $g \in \partial f$ for some convex $f : X \rightarrow \mathbf{R}$. Then

$$f(x') \geq f(x) + g(x) \cdot (x' - x) = f(x) + w(x, x') \quad \forall x, x' \in X,$$

so f is a potential for (X, w) . By Theorem 9, it follows that w satisfies (5), and hence g is cyclically monotone. \blacksquare

It remains to prove Theorem 9. The proof in Vohra (2011) uses the max-flow min-cut theorem; we'll give a direct proof.

Proof of Theorem 9. Suppose that f is a potential for (X, w) . Then for any finite cycle $x_0, x_1, \dots, x_{n-1}, x_n, x_0$,

$$\begin{aligned} f(x_1) &\geq f(x_0) + w(x_0, x_1) \\ f(x_2) &\geq f(x_1) + w(x_1, x_2) \\ &\vdots \\ f(x_n) &\geq f(x_{n-1}) + w(x_{n-1}, x_n) \\ f(x_0) &\geq f(x_n) + w(x_n, x_0). \end{aligned}$$

Summing yields (5).

Suppose that (5) holds for any finite cycle $x_0, x_1, \dots, x_{n-1}, x_n, x_0$. Rearranging,

$$\sum_{i=1}^n w(x_{i-1}, x_i) \leq -w(x_n, x_0).$$

Fix an $\bar{x} \in X$. Then for any given $x \in X$, $-w(x, \bar{x}) \in \mathbf{R}$ is an upper bound for $\sum_{i=1}^n w(x_{i-1}, x_i)$ across all finite cycles $x_0, x_1, \dots, x_{n-1}, x_n, x_0$ with $x_0 = \bar{x}$ and $x_n = x$. Hence the supremum exists (in \mathbf{R}) for each x . Let's call it $f(x)$:

$$f(x) := \sup \left\{ \sum_{i=1}^n w(x_{i-1}, x_i) : \right. \\ \left. x_0, x_1, \dots, x_{n-1}, x_n, x_0 \text{ a finite cycle with } x_0 = \bar{x} \text{ and } x_n = x \right\}.$$

f maps X into \mathbf{R} . Moreover, for any $x \in X$, we have

$$f(x) \geq w(\bar{x}, x_1) + w(x_1, x_2) + \cdots + w(x_{n-2}, x_{n-1}) + w(x_{n-1}, x)$$

for any finite cycle $x_0, x_1, \dots, x_{n-1}, x_n, x_0$ with $x_0 = \bar{x}$ and $x_n = x$, so if in particular $x_{n-1} = x'$ then

$$f(x) \geq [w(\bar{x}, x_1) + w(x_1, x_2) + \cdots + w(x_{n-2}, x')] + w(x', x).$$

Taking the supremum over the $[\cdot]$ term then yields

$$f(x) \geq f(x') + w(x', x).$$

Since $x, x' \in X$ were arbitrary, we've shown that f is a potential. ■

We saw in Theorem 8 that by replacing ‘increasing’ with ‘cyclically monotone’, we were able to recover most of the fundamental theorem of calculus from the one-dimensional case. But a more natural notion of monotonicity in higher dimensions is surely the following:

Definition 7. Let $X \subseteq \mathbf{R}^m$ be compact and convex. $g : X \rightarrow \mathbf{R}^m$ is two-cycle monotone iff

$$g(x) \cdot (x' - x) + g(x') \cdot (x - x') \leq 0 \quad \text{for every } x, x' \in X.$$

Besides being more natural, two-cycle monotonicity is far more tractable than cyclical monotonicity, since it is a ‘one-dimensional’ condition.

Cyclical monotonicity obviously implies two-cycle monotonicity. The converse is false in general; see Rockafellar (1970, p. 240) for a counterexample. But the converse is true in the one-dimensional case ($m = 1$). Are other interesting cases in which two-cycle monotonicity is sufficient for cyclical monotonicity? Here's one:

Theorem 10 (Saks and Yu (2005)). Let $X \subseteq \mathbf{R}^m$ be closed and convex, and let $g : X \rightarrow \mathbf{R}^m$ be such that $g(X)$ is finite. Then if g is two-cycle monotone, it is cyclically monotone.

Three quite different proofs can be found in Saks and Yu (2005), Vohra (2011, pp. 46–8) and Ashlagi, Braverman, Hassidim and Monderer (2010).¹⁷ The latter paper also provides a converse, identifying domains X on which two-cycle monotonicity implies cyclical monotonicity.

¹⁷A sketch proof based on Ashlagi et al. (2010) was given in class.

It might be instructive to note that the notions of monotonicity in this section resemble concepts from demand theory. In particular, cyclical monotonicity is a weakened form of SARP, two-cycle monotonicity is a weakened form of WARP, and rearranging two-cycle monotonicity as

$$[g(x') - g(x)] \cdot [x' - x] \geq 0$$

shows that it is equivalent to the law of demand. For quasi-linear preferences, SARP and WARP can in fact be replaced with the weaker conditions of cyclical monotonicity and two-cycle monotonicity. Our implementability result is closely analogous to Richter's theorem (demand is rationalisable iff it satisfies SARP). In general, SARP is stronger than WARP. In certain cases, such as when there are two goods (one-dimensional case), SARP and WARP are equivalent.

1.3.2 Environment

Consider a variant of the single-agent setting from section 1.1.3. Allocations are now $\gamma \in \Gamma \subseteq \mathbf{R}^m$ rather than $\in [0, 1]$, and types are correspondingly $t \in \mathcal{T} \subseteq \mathbf{R}^m$, where \mathcal{T} is compact and convex. Type t 's valuation of allocation γ is $\gamma \cdot t$.¹⁸ Otherwise the environment is the same as in section 1.1.3.

A direct mechanism is (g, p) , where $g : \mathcal{T} \rightarrow \Gamma$ is the allocation rule and $p : \mathcal{T} \rightarrow \mathbf{R}$ the transfer rule. IC, IR and implementability are defined exactly as in section 1.1.2:

Definition 8. Let (g, p) be a direct mechanism.

- (1) (g, p) is incentive-compatible (IC) iff

$$g(t) \cdot t - p(t) \geq g(t') \cdot t - p(t') \quad \text{for every } t, t' \in \mathcal{T}.$$

- (2) (g, p) is individually rational (IR) iff $g(t) \cdot t - p(t) \geq 0$ for every $t \in \mathcal{T}$.

- (3) g is implementable iff there is a $p : \mathcal{T} \rightarrow \mathbf{R}$ s.t. (g, p) is IC.

A revelation principle still applies. Our goals are to characterise implementability, to establish revenue equivalence, and to find the optimal mechanism.

¹⁸This is actually not very restrictive, though it has more bite when \mathcal{T} is assumed to be convex. Some of the results in this section can be obtained without 'dot-product' valuations; see Vohra (2011, ch. 4).

1.3.3 Characterisation of IC

In this section, we extend the characterisation of incentive compatibility in section 1.1.3 to the multidimensional setting. Rochet's theorem shows that an allocation rule g is implementable iff it is increasing in an appropriate sense. Theorem 3 shows that, up to a constant, any IC mechanism implementing g must have the same payments and payoffs.

As in the one-dimensional case, the key step is to write implementability in terms of the subdifferential of a convex function. In fact, the connection is exactly the same as in the one-dimensional case, and the proof is unchanged.

Lemma 4. $g : \mathcal{T} \rightarrow \Gamma$ is implementable iff $g \in \partial\phi$ for some convex $\phi : \mathcal{T} \rightarrow \mathbf{R}$.

Proof. Suppose that g is implementable. Then there is $p : \mathcal{T} \rightarrow \mathbf{R}$ such that (g, p) is IC. Define $\phi : \mathcal{T} \rightarrow \mathbf{R}$ from g and p as

$$\phi(t) := \max_{t' \in \mathcal{T}} [g(t') \cdot t - p(t')] \quad \text{for each } t \in \mathcal{T}.$$

Since ϕ is the pointwise maximum of convex functions, it is convex. Moreover, for every $t, t' \in \mathcal{T}$,

$$\begin{aligned} \phi(t') &\geq g(t) \cdot t' - p(t) \\ &= [g(t) \cdot t - p(t)] + g(t) \cdot (t' - t) \\ &\geq \phi(t) + g(t) \cdot (t' - t), \end{aligned}$$

where the first inequality used the definition of ϕ , and the second used IC. So $g \in \partial\phi$ for ϕ convex.

Now suppose that $g \in \partial\phi$ for some $\phi : \mathcal{T} \rightarrow \mathbf{R}$ convex. Define $p : \mathcal{T} \rightarrow \mathbf{R}$ from g and ϕ by

$$p(t) := g(t) \cdot t - \phi(t) \quad \text{for each } t \in \mathcal{T}.$$

Then for every $t, t' \in \mathcal{T}$,

$$\begin{aligned} g(t) \cdot t - p(t) &= \phi(t) \\ &\geq \phi(t') + g(t') \cdot (t - t') \\ &= [g(t') \cdot t' - p(t')] + g(t') \cdot (t - t') \\ &= g(t') \cdot t - p(t'), \end{aligned}$$

where the first and second equalities used the definition of p , and the inequality used $g(t') \in \partial\phi(t')$. So (g, p) is IC, hence g is implementable. \blacksquare

Since a function is a subdifferential of a convex function iff it is cyclically monotone, we obtain the analog of the Spence–Mirrlees theorem (p. 9):

Theorem 11 (Rochet (1987)). g is implementable iff it is cyclically monotone.

Proof. Immediate from Lemma 4 and part (1) of Theorem 8 (p. 34). ■

Corollary 13. If $g(\mathcal{T})$ is finite, then g is implementable iff it is two-cycle monotone.

Proof. Immediate from Rochet’s theorem and the Saks–Yu theorem (p. 37). ■

As above, let $\phi^{g,p}(t)$ denote the payoff of type t from the mechanism (g, p) . As shown in the proof of Lemma 4, $\phi^{g,p}$ is convex, and $g \in \partial\phi^{g,p}$ if (g, p) is IC. Since a convex function is determined up to a constant by its subdifferential, we have an envelope formula:

Theorem 12. If (g, p) is IC then

$$\phi^{g,p}(t') - \phi^{g,p}(t) = \int_0^1 g(t + \lambda(t' - t)) \cdot (t' - t) d\lambda \quad \text{for every } t, t' \in \mathcal{T}.$$

Proof. Immediate from Lemma 4 and part (2) of Theorem 8. ■

It follows that

Corollary 14 (payoff & revenue equivalence). If (g, p) and (g, p') are IC then

$$\phi^{g,p} - \phi^{g,p'} = p' - p = k$$

for some $k \in \mathbf{R}$.

1.3.4 Optimal mechanism

In section 1.1.4 (p. 10), we found the optimal mechanism for the a single agent with a one-dimensional type using a simple convexity argument. Let’s try to mimic that line of reasoning.

Theorem 12 implies that an IC allocation rule g pins down payments up to a constant. To maximise revenue, we choose the constant so that IR just holds. At least one type’s IR constraint will bind; write \underline{t} for one of

these types.¹⁹ Using the envelope formula in Theorem 12, we can then write revenue as a functional $R : \Gamma^{\mathcal{T}} \rightarrow \mathbf{R}$ of the allocation rule alone:

$$R(g) = \mathbf{E} \left(g(T) \cdot T - \int_0^1 g(\underline{t} + \lambda(T - \underline{t})) \cdot (T - \underline{t}) d\lambda \right).$$

Let \mathcal{M} denote the set of cyclically monotone maps $\mathcal{T} \rightarrow \Gamma$. \mathcal{M} is convex and compact. R is obviously linear. So as in the one-dimensional-types setting, maximising revenue amounts to maximising a linear functional R on a convex and compact space \mathcal{M} .

It follows that any optimal mechanism is an extreme point of \mathcal{M} . But it is very difficult to characterise the extreme points of \mathcal{M} , never mind to figure out which extreme points are optima. For example, it is not the case that all extreme points of \mathcal{M} are deterministic mechanisms, so we cannot rule out that randomisation is strictly optimal.

As an illustration, consider selling k objects to a single buyer: lotteries over allocations $\Gamma = \Delta(\{0, 1\}^k)$ and valuations $\mathcal{T} \subseteq \mathbf{R}^{2^k}$. (This setting is often called ‘combinatorial auctions’.) Even for the case of two objects with independently distributed valuations, almost nothing is known about the optimum! Results are available only under extreme assumptions, e.g. when valuations are drawn from a two-point distribution.

Hart and Reny (2015) provide examples of some of the strange properties that optimal mechanisms for selling k objects may have. One is ‘nonmonotonicity’: an increase (in the sense of first-order stochastic dominance) in valuations can decrease maximal revenue. Another is that it may be strictly optimal to use a random mechanism.

Not to say that there has been no progress. Border’s (1991) theorem can be used reformulate the revenue-maximisation problem in a way that is tractable in some special cases—see section 6.2.3 in Vohra (2011). There are nice results on revenue-maximising combinatorial auctions with gross-substitutes preferences. And much more besides.

¹⁹If \mathcal{T} is a lattice then \underline{t} is (uniquely) the pointwise minimum of \mathcal{T} .

1.4 Several agents, multidimensional types

1.4.1 Environment

We study the setting from section 1.3, but now with n agents. In particular, allocations are $\gamma \in \Gamma \subseteq \mathbf{R}^m$, types are $t_i \in \mathcal{T}_i \subseteq \mathbf{R}^m$ where \mathcal{T}_i is convex, and valuations are $\gamma \cdot t_i$. Write $\mathcal{T} = \prod_{i=1}^n \mathcal{T}_i$. A direct mechanism is (g, p) , where $g : \mathcal{T} \rightarrow \Gamma$ is the allocation rule and $p : \mathcal{T} \rightarrow \mathbf{R}^n$ the transfer rule.²⁰

1.4.2 Strategy-proof implementation

We're going to concern ourselves exclusively with strategy-proof implementation. The definitions are unchanged from section 1.2.4:

Definition 9. Let (g, p) be a direct mechanism.

- (1) (g, p) is dominant-strategies incentive-compatible (DSIC) iff for each i and each $t_{-i} \in \mathcal{T}_{-i}$,

$$g_i(t_i, t_{-i}) \cdot t_i - p_i(t_i, t_{-i}) \geq g_i(t'_i, t_{-i}) \cdot t_i - p_i(t'_i, t_{-i}) \quad \forall t_i, t'_i \in \mathcal{T}_i.$$

- (2) g is implementable in dominant strategies iff there is a $p : \mathcal{T} \rightarrow \mathbf{R}^n$ such that (g, p) is DSIC.

A revelation principle allows us to restrict attention to DSIC (and IR) direct revelation mechanisms.

As with one-dimensional types, it is trivial to extend the characterisation of IC for a single agent to a characterisation of DSIC for multiple agents because of the following observation.

Observation 3. A multi-agent mechanism (g, p) is DSIC iff each of the single-agent mechanisms $(g_i(\cdot, t_{-i}), p_i(\cdot, t_{-i}))$ are IC in the sense of section 1.3.2.

Corollary 15. g is implementable in dominant strategies iff each $g_i(\cdot, t_{-i})$ is cyclically monotone.

Proof. Immediate from Rochet's theorem (p. 40). ■

For a mechanism (g, p) , let $\phi_i^{g,p} : \mathcal{T} \rightarrow \mathbf{R}$ be given by

$$\phi_i^{g,p}(t_i, t_{-i}) := \max_{t'_i \in \mathcal{T}_i} [g_i(t'_i, t_{-i}) \cdot t_i - p_i(t'_i, t_{-i})] \quad \text{for each } (t_i, t_{-i}) \in \mathcal{T}.$$

²⁰We don't need \mathcal{T} to be compact because the main result of this section, Roberts's theorem, relies on unbounded valuations. We don't impose that types are drawn independently because we'll focus on strategy-proof implementation.

$\phi_i^{g,p}(t_i, t_{-i})$ is the (ex-post) payoff that type t_i of agent i earns in the mechanism when the others report t_{-i} .

Corollary 16. If (g, p) is DSIC then for each i and $t_{-i} \in \mathcal{T}_{-i}$,

$$\phi_i^{g,p}(t'_i, t_{-i}) - \phi_i^{g,p}(t_i, t_{-i}) = \int_0^1 g(t_i + \lambda(t'_i - t_i), t_{-i}) \cdot (t'_i - t_i) d\lambda$$

for every $t_i, t'_i \in \mathcal{T}_i$.

Proof. Immediate from Theorem 12 (p. 40). ■

Therefore

Corollary 17 (payoff & revenue equivalence). If (g, p) and (g, p') are DSIC then for each i and $t_{-i} \in \mathcal{T}_{-i}$,

$$\phi_i^{g,p}(\cdot, t_{-i}) - \phi_i^{g,p'}(\cdot, t_{-i}) = p'_i(\cdot, t_{-i}) - p_i(\cdot, t_{-i}) = k_{i,t_{-i}}$$

for some $k_{i,t_{-i}} \in \mathbf{R}$.

1.4.3 Roberts's theorem

In this section, we specialise to the case of a finite number of outcomes and unbounded valuations. In particular, consider a finite set A of outcomes, and let valuations lie in $\mathcal{T}_i = \mathbf{R}_+^A$. To fit this into the model above, let $\Gamma = \{e_a\}_{a \in A}$ be the basis vectors,²¹ so that type t_i 's valuation of outcome a is $e_a \cdot t_i =: t_i^a$. Also write $t^a := (t_1^a, \dots, t_n^a)$.

Observe that in this setting, \mathcal{T} is closed and convex, and $g(\mathcal{T})$ is finite for any g . Two-cycle monotonicity is therefore equivalent to cyclical monotonicity by the Saks–Yu theorem (p. 37), so that implementability is equivalent to two-cycle monotonicity.

An allocation $g : \mathcal{T} \rightarrow \Gamma$ is efficient iff

$$g(t) \in \arg \max_{a \in A} \sum_{i=1}^n t_i^a \quad \forall t \in \mathcal{T}.$$

Define $\phi : \mathcal{T} \rightarrow \mathbf{R}$ by

$$\phi(t_i, t_{-i}) := \max_{a \in A} \sum_{i=1}^n t_i^a = \max_{a \in A} \left(e_a \cdot t_i + \sum_{j \neq i} t_j^a \right) \quad \forall (t_i, t_{-i}) \in \mathcal{T}.$$

²¹Note that this rules out random mechanisms.

Then g is efficient iff $g(\cdot, t_{-i}) \in \partial\phi(\cdot, t_{-i})$ for every $t_{-i} \in \mathcal{T}_{-i}$, and ϕ is convex since it is the pointwise maximum of affine functions. Hence $g(\cdot, t_{-i})$ is cyclically monotone by Theorem 8 (p. 34), so implementable in dominant strategies by Corollary 15. This shouldn't be news: we already knew that any efficient g can be implemented in dominant strategies by a VCG mechanism.

There are other allocations that are obviously implementable in dominant strategies, e.g. constant and dictatorial allocations. But there aren't many: Roberts's theorem says that only affine maximisers are implementable.

Theorem 13 (Roberts (1979)). Suppose that $3 \leq |A| < \infty$, $\Gamma = \{e_a\}_{a \in A}$ and $\mathcal{T}_i = \mathbf{R}_+^{|A|}$ for each i . Then g is onto and implementable in dominant strategies and iff

$$g(t) \in \arg \max_{a \in A} \left[k(a) + \sum_{i=1}^n w(i)(e_a \cdot t_i) \right] \quad \text{for each } t \in \mathcal{T}$$

for some $k : A \rightarrow \mathbf{R}$ and $w : \{1, \dots, n\} \rightarrow [0, 1]$.

Remark 3. (Actually several remarks.)

- (1) This statement is from Vohra (2011, sec. 4.2.4). Roberts's (1979) original result actually had $\mathcal{T}_i = \mathbf{R}^{|A|}$.
- (2) Both $|A| \geq 3$ and the 'complete domain' condition that \mathcal{T}_i be all of $\mathbf{R}_+^{|A|}$ (or $\mathbf{R}^{|A|}$) are essential; see Vohra (2011, pp. 58–9) for counterexamples.
- (3) It is also essential that g be onto. But this is really without loss, since if g fails to be onto we can replace Γ by $g(\mathcal{T})$. (Provided, of course, that $|g(\mathcal{T})| \geq 3$.)
- (4) Vohra (2011, p. 54) says that it's easy to extend the theorem to infinite A .

The argument that we used to show that any efficient allocation is implementable in dominant strategies can obviously be adapted to show that any affine maximiser is implementable in dominant strategies, and it's obvious that any affine maximiser is onto. Proving the converse turns out to be rather a lot of work; see Vohra (2011, sec. 4.2.4).²²

Other characterisations of implementability in the spirit of Roberts's theorem include Bikhchandani et al. (2006) and Chung and Ely (2006). The former is for strategy-proof implementation, and the latter is for ex-post implementation.

²²A proof was given in class, but I omit it because it was long and hard. It was essentially the proof in Vohra (2011), but it simplified one step by using two-cycle monotonicity instead of cyclical monotonicity.

1.5 Several agents, correlated types

1.5.1 Environment

So far, we assumed that types are independent. In this section, we relax that assumption.

Consider a generalisation of the environment in section 1.4.1. Allocations are $\gamma \in \Gamma$, types are $t_i \in \mathcal{T}_i$, valuations are given by $v_i : \Gamma \times \mathcal{T}_i \rightarrow \mathbf{R}$, and preferences are quasi-linear in valuations and transfers.²³ Write $\mu_i(\cdot|t_i) \in \Delta(\mathcal{T}_{-i})$ for type t_i 's interim belief about the other agents' types, and let \mathbf{E}_{t_i} be shorthand for integration w.r.t. $\mu_i(\cdot|t_i)$.

A direct mechanism is (g, p) , where $g : \mathcal{T} \rightarrow \Gamma$ is the allocation rule and $p : \mathcal{T} \rightarrow \mathbf{R}^n$ the transfer rule. We're going to study Bayesian implementation in this setting.

Definition 10. Let (g, p) be a direct mechanism.

- (1) (g, p) is Bayes incentive-compatible (BIC) iff for each i ,

$$\begin{aligned} \mathbf{E}_{t_i} (v_i(g(t_i, T_{-i}), t_i) - p_i(t_i, T_{-i})) \\ \geq \mathbf{E}_{t_i} (v_i(g(t'_i, T_{-i}), t_i) - p_i(t'_i, T_{-i})) \quad \forall t_i, t'_i \in \mathcal{T}_i. \end{aligned}$$

- (2) (g, p) is interim individually rational (IIR) iff for each i ,

$$\mathbf{E}_{t_i} (v_i(g(t_i, T_{-i}), t_i) - p_i(t_i, T_{-i})) \geq 0 \quad \forall t_i \in \mathcal{T}_i.$$

- (3) g is implementable in Bayes–Nash iff there is a $p : \mathcal{T} \rightarrow \mathbf{R}^n$ such that (g, p) is BIC.

The revelation principle continues to hold with correlated types, allowing us to restrict attention to BIC (and IIR) direct mechanisms.

1.5.2 Characterisation of BIC (anything goes)

For the case of independent types, Rochet's theorem told us that the implementable allocation rules are precisely the cyclically monotone ones. This is a nontrivial condition: a great many allocation rules are not cyclically monotone, and so cannot be implemented. By contrast, the following theorem shows that when types are correlated (in a particular 'full rank' manner), every allocation rule is implementable.

²³In section 1.4.1, we specialised to $\Gamma \subseteq \mathbf{R}^m$, $\mathcal{T}_i \subseteq \mathbf{R}^m$ convex and $v_i(\gamma, t_i) = \gamma \cdot t_i$.

Theorem 14 (Crémer and McLean (1988), McAfee and Reny (1992)). Assume that for every i and $t_i \in \mathcal{T}_i$,

$$\mu_i(\cdot|t_i) \notin \text{co} \{ \mu_i(\cdot|t'_i) : t'_i \in \mathcal{T}_i \setminus \{t_i\} \}.$$

Then for every direct mechanism (g, p) , there is a $p' : \mathcal{T} \rightarrow \mathbf{R}^n$ such that (g, p') is IC, and for each i

$$\mathbf{E}_{t_i}(p_i(t_i, T_{-i})) = \mathbf{E}_{t_i}(p'_i(t_i, T_{-i})) \quad \forall t_i \in \mathcal{T}_i.$$

An obvious corollary is that if (g, p) is IIR, then (g, p') is IIR. Another obvious corollary is that an auctioneer can extract the full surplus from the agents. For whatever reason, the literature has tended to focus on this latter corollary rather than on the ‘anything goes’ characterisation of BIC.

The result as stated is from McAfee and Reny (1992), and so is the elegant method of proof below. These authors actually allow for infinite-dimensional type spaces.²⁴ Crémer and McLean (1988) proved essentially the same result in a more restrictive auction environment akin to the setting in section 1.4.1 above.²⁵

We can think of the condition on interim beliefs as a ‘spanning’ condition: its spirit is that different types’ beliefs are linearly independent. The condition is actually a lot weaker than that since the convex hull of a collection of vectors is an (often small) subset of their linear span. In fact, the spanning condition is so weak that it holds generically (when type spaces are finite-dimensional).

The idea of the proof is as follows. Types with linearly independent beliefs about the other players’ types have different preferences over gambles over the other players’ types. This means that we can construct side bets that are attractive to some types but not to others. By making these side bets very large, we ensure that no type wishes to lie, since lying means accepting a gamble whose (subjective) expected value is very large and negative.

Proof. Fix i and $t_i \in \mathcal{T}_i$. By the spanning condition and a (strict) separating hyperplane theorem, there exists $B : \mathcal{T}_{-i} \rightarrow \mathbf{R}$ such that $\mathbf{E}_{t_i}(B) = 0$ and $\mathbf{E}_{t'_i}(B) < 0$ for every $t'_i \in \mathcal{T}_i \setminus \{t_i\}$. B is our side bet.

²⁴Extending the result to infinite-dimensional type spaces is nontrivial, but it doesn’t change the economics, so we won’t bother.

²⁵As usual, Myerson got here first: in Myerson (1981, sec. 7), he gives an example of how an auctioneer can exploit correlation to extract the full surplus from the bidders in an auction setting. His mechanism involves offering the agents side bets about the other players’ types, making use of the fact that different types have different preferences over such bets. Our proof below constructs exactly such a mechanism.

Choose a large number $M \in \mathbf{R}$; in particular, big enough that

$$\begin{aligned} \mathbf{E}_{t'_i} (v_i(g(t'_i, T_{-i}), t'_i) - p_i(t'_i, T_{-i})) \\ \geq \mathbf{E}_{t'_i} (v_i(g(t_i, T_{-i}), t'_i) - p_i(t_i, T_{-i})) + M\mathbf{E}_{t'_i} (B) \quad \forall t'_i \in \mathcal{T}_i. \end{aligned} \quad (6)$$

Let the new payment rule satisfy

$$p'_i(t_i, \cdot) := p_i(t_i, \cdot) - MB.$$

Clearly $\mathbf{E}_{t_i} (p'_i(t_i, T_{-i})) = \mathbf{E}_{t_i} (p_i(t_i, T_{-i}))$ since $\mathbf{E}_{t_i} (B) = 0$.

Having done this for every i and $t_i \in \mathcal{T}_i$, we have a new payment rule $p' : \mathcal{T} \rightarrow \mathbf{R}$ with the same interim expected payments as p . Since we required (6) to hold for each i and t_i , (g, p') is BIC. \blacksquare

1.5.3 Eliciting beliefs

Consider a seemingly quite different problem. An agent has a belief about $x \in X$, and has risk-neutral preferences over monetary lotteries. A principal wishes to elicit the agent's belief, and her only tool is to offer the agent x -contingent lotteries. Can she do it? (Yes.)

Let X be finite and write $\Delta(X)$ for the set of all probability mass functions on X . A mechanism is a map $t : \Delta(X) \rightarrow \mathbf{R}^X$.²⁶ t is incentive-compatible (IC) iff

$$\sum_{x \in X} p(x)t(p(x)) \geq \sum_{x \in X} p(x)t(p'(x)) \quad \forall p, p' \in \Delta(X).$$

This is clearly a species of linear mechanism design problem. The argument used to establish Lemma 1 all the way back in section 1.1.3 (p. 8) applies here (*mutatis mutandis*), so t is incentive-compatible iff $t \in \partial\phi$ for some convex $\phi : \Delta(X) \rightarrow \mathbf{R}$.

One particular convex function is the entropy function

$$\phi(p) = \sum_{x \in X} p(x) \ln(p(x)).$$

The (unique) associated IC mechanism is $t(p') = \ln(p')$. In this case we can prove IC directly using Jensen's inequality:

$$\sum_{x \in X} p(x) \ln \left(\frac{p(x)}{p'(x)} \right) \leq \ln \left(\sum_{x \in X} p(x) \frac{p(x)}{p'(x)} \right) = \ln(1) = 0.²⁷$$

²⁶To extend this to general X , fix some reference measure and work with the densities of beliefs (measures) w.r.t. this reference measure. A mechanism maps densities (reported beliefs) into suitably integrable random variables.

²⁷The object on the LHS is called the Kullback–Leibler divergence between p and p' . The (trivial) result that it is nonpositive is called ‘the information inequality’ by some.

This is one way of eliciting the agent’s belief.

Since there are a lot of convex functions $\Delta(X) \rightarrow \mathbf{R}$, there are a lot of other IC mechanisms. The moral is that eliciting beliefs is easy. More: gambles can clearly be chosen to have prior expectation zero, so eliciting beliefs is genuinely for free.

There is a more general theory of belief elicitation using gambles. (Mainly in the statistics literature, I think.) Mechanisms are usually called ‘scoring rules’ in this context.

Aside We assumed risk-neutrality only for ease of exposition, but that is wlog for two reasons. First, if our agent has a von Neumann–Morgenstern utility function $u : \mathbf{R} \rightarrow \mathbf{R}$ that is onto then we can reinterpret all monetary lotteries above as utility lotteries (since any utility lottery can then be achieved via some monetary lottery). Secondly (and more profoundly), since even risk-neutral expected-utility maximisers’ preferences are linear in the probabilities, we can fix a money (or utility) prize and vary the winning probability. In particular, let a mechanism t take a reported belief p' and spit out a probability $t(p')$ of winning the fixed prize, so that an agent who reports p' gets the prize with probability $t(p')$.

1.5.4 More on the ‘anything goes’ result

Theorem 14 is a consequence of the fact that belief elicitation is easy. Regardless of how types are distributed, the principal can trivially elicit i ’s belief. When the spanning condition holds, she can infer i ’s type (preferences). So the information asymmetry is gone: the principal can elicit the type of every agent for free. It obviously follows that she can implement whatever allocation rule she wants by harshly punishing agents when they lie.²⁸

Everyone agrees that Theorem 14 is troubling; after all, we think that types are often at least slightly correlated, and yet we do not observe full-rent-extracting auctions using huge side bets. There is a fairly large cottage industry devoted to breaking the ‘anything goes’ result. One obvious solution is to impose ex post limited liability, i.e. that an agent can walk away if the mechanism asks her for a huge transfer. Another is to assume risk aversion and individual rationality, since risk-averse agents will not participate if they have to accept huge side bets.

The spanning condition obviously implies the pairwise restriction that no two types have the same beliefs. The contrapositive of this pairwise property

²⁸This observation might lead us to feel that Theorem 14 is not very interesting.

is that ‘beliefs determine preferences’ (BDP). Neeman (2004) showed by example that when BDP (and hence spanning) fails, IC becomes nontrivial again (in particular, he shows that full rent extraction may not be possible). Most of the literature on breaking the ‘anything goes’ result has attacked the BDP property.

Perhaps the most popular escape route has been to argue that BDP is non-generic.²⁹ A well-known ‘non-genericity’ paper is Heifetz and Neeman (2006). Bergemann and Morris (2005) have a simpler (and earlier) non-genericity result: the set of non-BDP types is dense in the product topology on the universal type space.³⁰ The ‘genericity’ papers are arguably just as compelling, though.

Eran has the more pessimistic view that a framework that gives us a result as crazy as Theorem 14 theorem must be broken. He seems to think that something quite radically different is needed. One way of trying to address this concern is to require incentive compatibility to hold on a larger subset of the universal type space. That is the course pursued by Bergemann and Morris in their series of papers on ‘robust’ mechanism design, starting with Bergemann and Morris (2005).

²⁹It is obviously generic when type spaces are finite-dimensional, but genericity in infinite-dimensional spaces is delicate. There are two main kinds of (non)genericity: topological (e.g. nowhere denseness) and measure-theoretic (e.g. shyness). In finite-dimensional spaces, it is hard to argue with the usual Euclidean topology and Lebesgue measure, and in fact topological and measure-theoretic notions of genericity (mostly) coincide. In infinite-dimensional spaces, neither the choice of topology nor the choice of measure are obvious, and topological and measure-theoretic notions of genericity generally fail to coincide. Eran’s reaction to this is that we shouldn’t talk about genericity in infinite-dimensional spaces; he contends that economists are the only people who insist on doing so.

³⁰They use the terminology ‘payoff type spaces’ rather than ‘BDP’.

2 Supermodular games

Consider the Diamond (1982) search model. There are n players, each choosing her effort level $x_i \in \mathbf{R}_+$. i 's payoff is

$$u_i(x_i, x_{-i}) = \alpha g \left(\frac{1}{n-1} \sum_{j \neq i} x_j \right) x_i - \frac{1}{2} x_i^2,$$

where $g : \mathbf{R} \rightarrow \mathbf{R}$ is strictly increasing and smooth with $g(0) > 0$, and $\alpha \in \mathbf{R}_{++}$ is a parameter. Restricting attention to symmetric equilibria, i 's best response is $\text{BR}_i^\alpha(x) = \alpha g(x)$.

The following are (immediate) properties of symmetric Nash equilibria.

- (1) They are all pure.
- (2) There are best and worst equilibria (in the Pareto order).
- (3) There are largest and smallest equilibria (they coincide with the best and worst ones).
- (4) The largest and smallest equilibria are increasing in α , though intermediate ones need not be.

Supermodular games, introduced by Milgrom and Roberts (1990) and Vives (1990), are a large-ish class of games to which these four properties generalise. The mathematical tools are mainly due to Don Topkis (especially his 1978). The (very dense) standard reference for this material is Topkis (1998).

2.1 Lattices

Definition 11. Let L be a set, and let \succsim be a binary relation on L . \succsim is a partial order on L iff it is reflexive, transitive and antisymmetric. (L, \succsim) is then called a poset (partially ordered set).

Example 1. Let \geq be the binary relation on \mathbf{R}^n given given by $x \geq y$ for iff $x_i \geq y_i$ for each $i = 1, \dots, n$. \geq is a partial order, so (\mathbf{R}^n, \geq) is a poset.

Fix a poset (L, \succsim) . If $x \in L$ and $R \subseteq L$, x is called an upper bound of R in (L, \succsim) iff $x \succsim y$ for every $y \in R$; similarly for lower bounds. The least upper bound of R in (L, \succsim) , if it exists, is denoted $\sup_{(L, \succsim)} R$, and similarly $\inf_{(L, \succsim)} R$ denotes the greatest lower bound. Figure 1 illustrates.

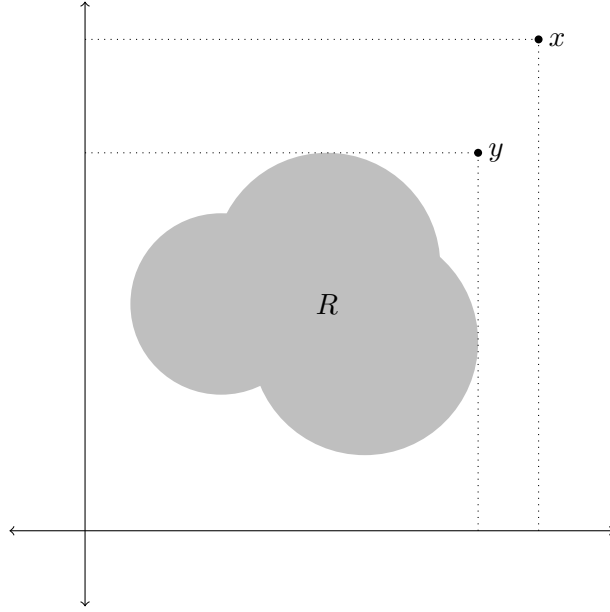


Figure 1 – The poset (\mathbf{R}^2, \geq) and a subset $R \subseteq \mathbf{R}^2$. x is an upper bound of R in (\mathbf{R}^2, \geq) , and y is the least upper bound (supremum) of R in (\mathbf{R}^2, \geq) .

Definition 12. A poset (L, \succsim) is a lattice iff

$$\sup_{(L, \succsim)} \{x, y\} \quad \text{and} \quad \inf_{(L, \succsim)} \{x, y\}$$

exist for every pair $x, y \in L$.

We will sometimes be sloppy and refer to a set L as a poset or lattice, suppressing the order (as we often do with topologies, metrics, etc.)

Definition 13. A lattice (L, \succsim) is complete iff

$$\sup_{(L, \succsim)} R \quad \text{and} \quad \inf_{(L, \succsim)} R$$

exist for every $R \subseteq L$.³¹

Example 2. The posets (\mathbf{R}^n, \geq) and (\mathbf{N}^n, \geq) are complete lattices. (\mathbf{Q}^n, \geq) is also a lattice, but it is not complete.

³¹Some authors use ‘complete’ to mean something weaker. We follow the convention from Topkis (1998).

Example 3. For any set Ω , $(2^\Omega, \supseteq)$ is a complete lattice.

We can generalise these examples as follows. Let $\{(L_i, \succsim_i)\}_{i \in I}$ be a collection of lattices, set $L := \prod_{i \in I} L_i$ and define \succsim on L by $x \succsim y$ iff $x_i \succsim_i y_i$ for each $i \in I$. (L, \succsim) thus defined is called a product lattice. As the name suggests, a product lattice is a lattice. Clearly (\mathbf{R}^n, \geq) is a product lattice. So is $(2^\Omega, \supseteq)$: it is the product of lattices $(\{0, 1\}, \geq)$.

Example 4. Let $\Delta([a, b])$ denote the set of Borel probability measures on $[a, b]$, and let \succsim_F be the first-order stochastic dominance order on $\Delta([a, b])$. $(\Delta([a, b]), \succsim_F)$ is a complete lattice.

We now consider the lattice structure of subsets of lattices. For a poset (L, \succsim) and a subset $R \subseteq L$, let $\succsim|_R$ be the restriction of \succsim to R . Clearly $(R, \succsim|_R)$ is also a poset. To save ourselves some pain, we will henceforth be sloppy and write simply (R, \succsim) , taking it as understood that \succsim is restricted to the appropriate domain.

Definition 14. Let (L, \succsim) be a lattice and $L' \subseteq L$. (L', \succsim) is a sublattice of (L, \succsim) iff for any $x, y \in L'$, $\sup_{(L, \succsim)}\{x, y\}$ and $\inf_{(L, \succsim)}\{x, y\}$ lie in L' .

Figure 2 gives some examples of sublattices and non-sublattices.

Obviously any sublattice is itself a lattice. But the converse is not true!

Example 5. Let $L := \{(1/2, 1/2), (2, 1), (1, 2), (3, 3)\}$, as drawn in Figure 3. (L, \geq) is not a sublattice of (\mathbf{R}^2, \geq) since for example

$$\sup_{(\mathbf{R}^2, \geq)} \{(2, 1), (1, 2)\} = (2, 2) \notin L.$$

It is a lattice, though; for example, $\sup_{(L, \geq)}\{(2, 1), (1, 2)\} = (3, 3)$.

For $1 \leq i < j \leq n$, let $\pi_{ij} : \mathbf{R}^n \rightarrow \mathbf{R}^2$ be the projection map

$$\pi_{ij}(x_1, \dots, x_i, \dots, x_j, \dots, x_n) = (x_i, x_j).$$

The following theorem says that for ‘Euclidean’ lattices (\mathbf{R}^n, \geq) , the sublattice property holds iff it holds for every pair of dimensions.

Theorem 15. Let (L, \geq) be a sublattice of (\mathbf{R}^n, \geq) , and take $x \in \mathbf{R}^n$. Then $x \in L$ iff $\pi_{ij}(x) \in \pi_{ij}(L)$ for all $1 \leq i < j \leq n$.

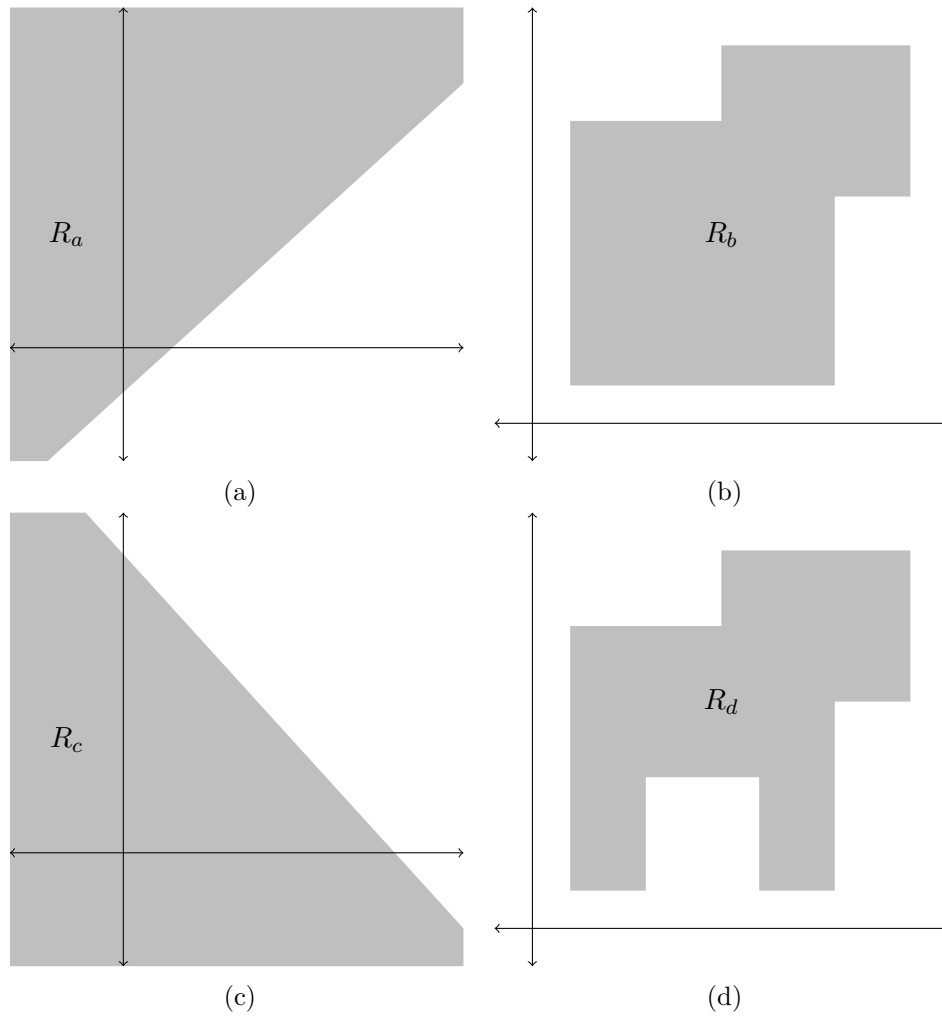


Figure 2 – The lattice (\mathbf{R}^2, \geq) and four subsets. (R_a, \geq) and (R_b, \geq) are sublattices of (\mathbf{R}^2, \geq) , whereas (R_c, \geq) and (R_d, \geq) are not. (In fact, the latter two are not even lattices in their own right.)




Figure 3 – The set $L = \{(\frac{1}{2}, \frac{1}{2}), (2, 1), (1, 2), (3, 3)\} \subseteq \mathbf{R}^2$. (L, \geq) is a lattice, but it is not a sublattice of (\mathbf{R}^2, \geq) .

Proof. Let (L, \geq) be a sublattice of (\mathbf{R}^n, \geq) , and take $x \in \mathbf{R}^n$. It's obvious that $x \in L$ implies $\pi_{ij}(x) \in \pi_{ij}(L)$ for any $1 \leq i < j \leq n$.

Let $\{v_{ij}\}_{1 \leq i < j \leq n}$ be vectors in \mathbf{R}^n such that $\pi_{ij}(v_{ij}) = \pi_{ij}(x) \in \pi_{ij}(L)$. Fix $i \in \{1, \dots, n\}$, and consider

$$V_i := \sup_{(\mathbf{R}^n, \geq)} \{v_{i1}, \dots, v_{i,j-1}, v_{i,j+1}, \dots, v_{in}\}.$$

The i th coordinate of V_i obviously equals x_i . The j th coordinate for $j \neq i$ is weakly greater than the j th coordinate of x . It follows that

$$\inf_{(\mathbf{R}^n, \geq)} \{V_1, \dots, V_n\} = x.$$

Since (L, \geq) is a sublattice of (\mathbf{R}^n, \geq) , $\inf_{(\mathbf{R}^n, \geq)} \{V_1, \dots, V_n\} \in L$; hence $x \in L$. ■

Definition 15. Let (L, \succsim) be a complete lattice and (L', \succsim) a sublattice of it. (L', \succsim) is subcomplete iff for every $R \subseteq L'$, $\sup_{(L', \succsim)} R$ and $\inf_{(L', \succsim)} R$ lie in L' .

A subcomplete sublattice is a sublattice. A subcomplete sublattice is also a complete lattice in its own right. But the converse is false!

Example 6. Let $L := \{0\} \cup (1/2, 1]$. (L, \geq) is a sublattice of (\mathbf{R}, \geq) , and is a complete lattice in its own right. (Convince yourself.) But it is not a subcomplete sublattice of (\mathbf{R}, \geq) since

$$\inf_{(\mathbf{R}, \geq)} (1/2, 1] = 1/2 \notin L.$$

Theorem 16. A sublattice (L, \geq) of (\mathbf{R}^n, \geq) is subcomplete iff L is compact.

The proof is straightforward, and perhaps a good exercise.

2.2 Tarski's fixed-point theorem

Definition 16. Let (L, \succsim) be a poset. A map $f : L \rightarrow L$ is increasing (or monotone, or \succsim -order-preserving) iff $x \succsim y$ implies $f(x) \succsim f(y)$.

We are interested in the set $\{x \in L : f(x) = x\}$ of fixed points of an increasing function f . We have the following beautiful theorem.

Theorem 17 (Tarski (1955)). If (L, \succsim) is a nonempty complete lattice and $f : L \rightarrow L$ is order-preserving, then the set of fixed points of f is a nonempty complete lattice with greatest and least elements

$$\sup_{(L, \succsim)} \{x \in L : x \preceq f(x)\} \quad \text{and} \quad \inf_{(L, \succsim)} \{x \in L : x \succeq f(x)\}.$$

Figure 4 illustrates the theorem and compares it with Brouwer's (1912) fixed-point theorem.

For L finite, it's obvious why this is true. Since L is a nonempty complete lattice, it has a smallest element $x_0 \in L$. For each $n \in \mathbf{N}$, define $x_n = f(x_{n-1})$. By monotonicity, $x_n = f(x_{n-1}) \succeq x_{n-1}$ for each $n \in \mathbf{N}$. By finiteness, strict improvements must eventually stop, leaving us with a fixed point.

The fact that the result holds for arbitrary L is what makes the theorem deep. We'll just prove the existence of a fixed point. (But the other steps are entirely straightforward.)

Proof of existence. Define

$$R := \{x \in L : x \succeq f(x)\}.$$

R is nonempty since $\sup_{(L, \succsim)} L \in R$. Define $x_\star := \inf_{(L, \succsim)} R$; x_\star is well-defined since R is nonempty and (L, \succsim) is complete.

Take any $x \in R$; then $x_\star \preceq x$, so $f(x_\star) \preceq f(x)$ by monotonicity. Since $f(x) \preceq x$ by definition of R , it follows that $f(x_\star) \preceq x$ for any $x \in R$, so that




Figure 4 – Two functions $f : L \rightarrow L$, where L is a compact and convex subset of \mathbf{R} and (hence) (L, \geq) is a complete lattice. Panel (a) illustrates Brouwer’s fixed point theorem, which requires continuity of f . Panel (b) depicts Tarski’s fixed-point theorem, which requires that f be increasing.

$f(x_*)$ is a lower bound of R . Since x_* is the greatest lower bound of R (by definition), $x_* \gtrsim f(x_*)$.

Hence $x_* \in R$, so that $x_* \gtrsim f(x_*)$. By monotonicity, it follows that $f(x_*) \gtrsim f(f(x_*))$, and hence $f(x_*) \in R$ by definition of R . Since x_* is a lower bound of R , it follows that $x_* \lesssim f(x_*)$.

Since \gtrsim is antisymmetric, we conclude that $x_* = f(x_*)$. ■

Tarski’s fixed-point theorem can be used for all kinds of things. Some familiar examples are as follows.

Example 7 (best-reply sets). Consider a game $(I, \{A_i, u_i\}_{i \in I})$ in normal form. Let $B_i(A'_{-i})$ denote the set of best responses $\in A_i$ to conjectures supported on $A'_{-i} \subseteq A_{-i}$. Write $A := \prod_{i \in I} A_i$ and $B(A) := (B_i(A_{-i}))_{i \in I}$. A product set $A' := \prod_{i \in I} A'_i \subseteq A$ has the best-reply property iff for any profile $(a_i)_{i \in I} \in A'$ and any player $i \in I$, a_i is a best reply to some conjecture supported on A'_{-i} . More succinctly, A' has the best-reply property iff $A' \subseteq B(A')$.

Let $\mathcal{P}(A)$ denote the product subsets of A . $(\mathcal{P}(A), \supseteq)$ is clearly a complete lattice, and $B : \mathcal{P}(A) \rightarrow \mathcal{P}(A)$ is clearly \supseteq -order-preserving. Hence by Tarski’s fixed-point theorem, B has a largest fixed point. It follows that there is a

largest best-reply set R (R is a best-reply set and contains every other best-reply set).

For $k \in \mathbf{N}$, let A^k be the (product) set of strategy profiles surviving k rounds of deletion of never-best-replies, so that $A^\infty := \bigcap_{k \in \mathbf{N}} A^k$ is the rationalisable set. Clearly $A^k \supseteq R$ for each $k \in \mathbf{N}$, hence $A^\infty \supseteq R$. Each A^k has the best-reply property by definition, and if action spaces are compact and payoffs continuous,³² then the best-reply property is preserved under countable intersection, so $A^\infty \subseteq B(A^\infty)$. It follows that $A^\infty = R$, i.e. the rationalisable set coincides with the maximal best-reply set.³³

Example 8 (repeated games). Consider a repeated game with imperfect monitoring, and let the set of players be I . As in Abreu, Pearce and Stacchetti (1990), let $B(W) \subseteq \mathbf{R}^{|I|}$ denote the set of payoff profiles generated by the payoff profiles $W \subseteq \mathbf{R}^{|I|}$. A set $W \subseteq \mathbf{R}^{|I|}$ is called self-generating iff $W \subseteq B(W)$.

Clearly $(2^{\mathbf{R}^n}, \supseteq)$ is a complete lattice and B is \supseteq -order-preserving. Hence by Tarski's theorem, B has a largest fixed point: there is a self-generating set W^* that contains all other self-generating sets. W^* is the set of all PPE payoff vectors.

Example 9 (Markov chains). Consider a Markov chain on $[a, b]$ with transition kernel $f : [a, b] \rightarrow \Delta([a, b])$.³⁴ The induced Markov operator $T : \Delta([a, b]) \rightarrow \Delta([a, b])$ is given by

$$(T\mu)(E) := \int_{[a,b]} f(E|x)\mu(dx) \quad \text{for every } E \subseteq [a, b] \text{ Borel.}$$

$\mu \in \Delta([a, b])$ is an invariant distribution of the Markov chain iff it is a fixed point of T .

Let's assume that $x \geq y$ implies $f(x) \succsim_F f(y)$, where \succsim_F is the first-order stochastic dominance order. This obviously implies that T is \succsim_F -order-preserving. We already know that $(\Delta([a, b]), \succsim_F)$ is a complete lattice. It follows by Tarski's theorem that the Markov chain has an invariant distribution, and further that its invariant distributions form a complete lattice with respect to first-order stochastic dominance.³⁵

³²Finite action spaces (with the discrete topology) are a special case.

³³Note that the existence, nonemptiness and properties of the largest best-reply set requires no assumptions at all on the game. By contrast, the nonemptiness and properties of the rationalisable set rely on topological structure.

³⁴ $\Delta([a, b])$ are the Borel measures on $[a, b]$, remember.

³⁵The other obvious way to guarantee the existence of an invariant distribution is to

Example 10 (matching). Let W and M be sets of heterosexual women and men, respectively. For simplicity, assume that $|W| = |M| = n$. Each woman $w \in W$ has a strict preference relation \succ_w on M ,³⁶ and similarly men $m \in M$ have preferences \succ_m over W .

A matching is a map $\mu : W \cup M \rightarrow W \cup M$ such that $\mu(W) \subseteq M$ and $\mu(M) \subseteq W$ (heteronormativity) and $\mu \circ \mu$ is the identity (consistency). A matching μ is stable iff there is no pair $(w, m) \in W \times M$ s.t. $m \succ_w \mu(w)$ and $w \succ_m \mu(m)$.

A prematching is a pair (f, g) of maps $f : W \rightarrow M \cup \{\emptyset\}$ and $g : M \rightarrow W \cup \{\emptyset\}$. (Unlike a matching, a prematching allows people to remain single, and doesn't impose consistency.) Let L be the set of all prematchings, and write $(f, g) \succeq (f', g')$ iff all women weakly prefer f to f' and all men weakly prefer g' to g . Surprisingly (at least to me), (L, \succeq) is a complete lattice. (The proof is not hard.)

For a prematching (f, g) , let Tf be the map $W \rightarrow M \cup \{\emptyset\}$ we get by letting each woman w grab her favourite man from among those who are either single or are willing to leave their present spouse under g for w , and similarly for Tg . Formally, $T : L \rightarrow L$ is given by

$$Tf(w) := \max_{\succ_w} \{m \in M : g(m) = \emptyset \text{ or } w \succeq_m g(m)\}$$

$$Tg(m) := \max_{\succ_m} \{w \in W : f(w) = \emptyset \text{ or } m \succeq_w f(w)\}.$$

An application of T to a prematching is essentially equivalent to one round of the Gale–Shapley (1962) algorithm. A moment's reflection reveals that T is \succeq -order-preserving.

It follows by Tarski's fixed-point theorem that the set R of fixed points of T is a nonempty complete lattice w.r.t. \succeq . Fixed points of T clearly correspond to stable matchings, so we've shown not only that stable matchings exist, but that there is a stable matching preferred by all women (viz. $\sup_{(L, \succeq)} R$) and a stable matching preferred by all men ($\inf_{(L, \succeq)} R$). The latter result is not obvious, nor even intuitive! (I would say.)

Since L is finite, starting at the smallest or largest prematching and iterating on T yields a stable matching in finite time. Starting the algorithm at the largest prematching will generate the women-favoured matching, and starting at the smallest prematching yields the men-preferred matching.³⁷

impose continuity and compactness assumptions, then to appeal to Brouwer's fixed-point theorem. The natural route would be to endow $\Delta([a, b])$ with the weak* topology (that gives us compactness for free) and to assume that f is continuous.

³⁶'Strict preference order' here just means 'total order'.

³⁷The lattice structure of stable matchings was first pointed out in print by Knuth (1976)

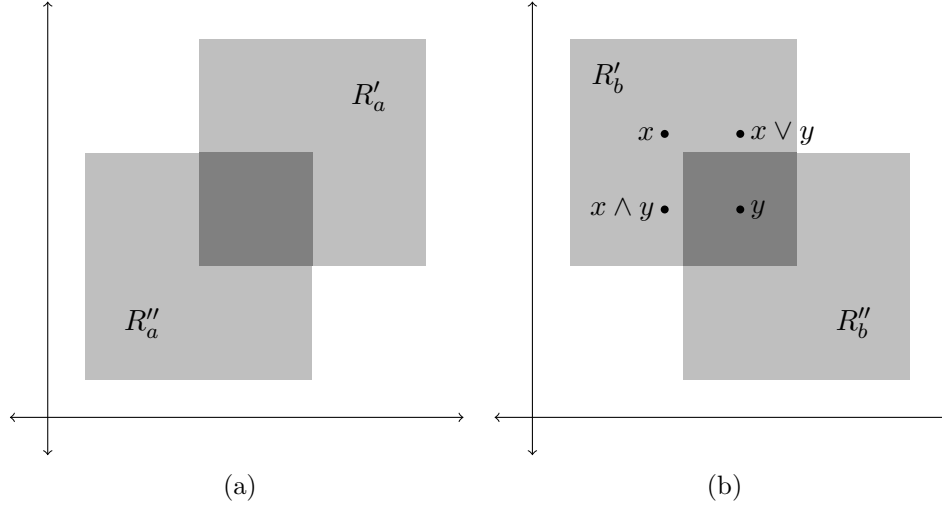


Figure 5 – The lattice (\mathbf{R}^2, \succeq) and some sublattices. In panel (a), $R'_a \sqsupseteq R''_a$. In panel (b), R'_b and R''_b are not ranked by the strong set order.

To study best-response correspondences and hence Nash equilibria, we will require an extension of Tarski's theorem to correspondences. To do that, we need an appropriate notion of monotonicity for correspondences, which in turn requires an order on sublattices.

Definition 17. Let $\mathcal{P}(L, \succeq)$ be the set of sublattices of a lattice (L, \succeq) . The strong set order induced by (L, \succeq) is the order \sqsupseteq on $\mathcal{P}(L, \succeq)$ such that $(L', \succeq) \sqsupseteq (L'', \succeq)$ iff for any $x' \in L'$ and $x'' \in L''$, $\sup_{(L, \succeq)}\{x', x''\} \in L'$ and $\inf_{(L, \succeq)}\{x', x''\} \in L''$.

Figure 5 illustrates the strong set order.

Definition 18. Let (L, \succeq) be a lattice, and let \sqsupseteq be the strong order it induces on $\mathcal{P}(L, \succeq)$. A correspondence $f : L \rightrightarrows L$ is increasing (in the induced order) iff it is sublattice-valued and $x \succeq y$ implies $f(x) \sqsupseteq f(y)$.³⁸

Theorem 18 (Zhou (1994)). If (L, \succeq) is a nonempty complete lattice and $f : L \rightrightarrows L$ is increasing in the induced order, then the set of fixed points of f is a nonempty complete lattice.

(that's the guy who invented \TeX !), but he attributes the discovery to John Conway.

³⁸If we view the correspondence f as a set-valued function, then this definition says that f must be sublattice-valued and \sqsupseteq -order-preserving.

2.3 Topkis's monotonicity theorem

Definition 19. Let (L, \succsim) and (L', \succsim') be posets. $f : L \times L' \rightarrow \mathbf{R}$ has increasing differences iff whenever $x \succsim x'$, $\delta(\cdot) := f(x, \cdot) - f(x', \cdot)$ is increasing in \succsim' .

Definition 20. Let (L, \succsim) be a lattice. $f : L \rightarrow \mathbf{R}$ is supermodular iff for any $x, x' \in L$,

$$f\left(\sup_{(L, \succsim)} \{x, x'\}\right) + f\left(\inf_{(L, \succsim)} \{x, x'\}\right) \geq f(x) + f(x').$$

For those who like pictures, Figure 6 illustrates what a supermodular function looks like.

Proposition 7. $f : \mathbf{R}^n \rightarrow \mathbf{R}$ is supermodular iff it has increasing differences in every pair of coordinates.

Note the similarity with our theorem about the structure of sublattices of (\mathbf{R}^n, \geq) (p. 52). The (straightforward) proof is left as an exercise.

The previous proposition obviously implies that for f twice continuously differentiable, supermodularity is equivalent to every off-diagonal element of the Hessian $\nabla^2 f$ being nonnegative. For comparison, convexity is equivalent to positive semidefiniteness of $\nabla^2 f$.

Now consider single-person decision problems. A problem is parameterised by $t \in T$, where (T, \succeq) is a poset. The decision-maker chooses an element of a sublattice $S(t)$ of a lattice (L, \succsim) to maximise $f(\cdot, t) : L \rightarrow \mathbf{R}$. Write

$$X(t) := \arg \max_{x \in S(t)} f(x, t).$$

With appropriate restrictions on the constraint correspondence $S : T \Rightarrow L$ and the objective $f : L \times T \rightarrow \mathbf{R}$, the argmax correspondence $X : T \Rightarrow L$ has a nice lattice structure and is increasing. In particular:

Lemma 5. Let (L, \succsim) be a lattice, and (S, \succsim) a sublattice of it. Assume that $f : L \rightarrow \mathbf{R}$ is supermodular. Then $\arg \max_{x \in S} f(x)$ is a sublattice of (L, \succsim) .

Theorem 19 (Topkis (1978)). Let (L, \succsim) be a lattice and (T, \succeq) a poset. Assume that $S : T \Rightarrow L$ is increasing in the induced order, that $f : L \times T \rightarrow \mathbf{R}$ has increasing differences, and that $f(\cdot, t)$ is supermodular for each $t \in T$. Then the argmax correspondence $X : T \Rightarrow L$ given by

$$X(t) := \arg \max_{x \in S(t)} f(x, t) \quad \forall t \in T$$




Figure 6 – The lattice (\mathbf{R}^2, \geq) and a supermodular function $f : \mathbf{R}^2 \rightarrow \mathbf{R}$. The symbols \vee and \wedge are shorthand for $\sup_{(\mathbf{R}^2, \geq)}$ and $\inf_{(\mathbf{R}^2, \geq)}$.

is increasing in the induced order.

Supermodularity can be weakened to quasi-supermodularity, and increasing differences can be weakened to single crossing (which should really be called single-crossing differences). There are various converses as well. See Topkis (1998) and Milgrom and Shannon (1994).

2.4 Pure-strategy Nash equilibria of supermodular games

A normal-form game is $(I, \{A_i, u_i\}_{i \in I})$ as usual.

Definition 21. $(I, \{A_i, u_i\}_{i \in I})$ is a supermodular game iff for each $i \in I$,

- (1) (A_i, \geq) is a subcomplete sublattice of (\mathbf{R}^{m_i}, \geq) for some $m_i \in \mathbf{N}$,
- (2) $u_i : A_i \times A_{-i} \rightarrow \mathbf{R}$ has increasing differences, and
- (3) $u_i(\cdot, a_{-i})$ is continuous and supermodular for each $a_{-i} \in A_{-i}$.

The set of pure-strategy Nash equilibria of a supermodular game has a nice structure:

Proposition 8. The set of pure-strategy Nash equilibria of a supermodular game is a nonempty complete lattice.

Proof. For any $i \in I$, the pure-strategy best-response correspondence $\text{BR}_i : A_{-i} \Rightarrow A_i$ is given by

$$\text{BR}_i(a_{-i}) := \arg \max_{a_i \in A_i} u_i(a_i, a_{-i}) \quad \forall a_{-i} \in A_{-i}.$$

The pure-strategy Nash correspondence $\text{BR} : A \Rightarrow A$ is defined by

$$\text{BR}(a) = \prod_{i \in I} \text{BR}_i(a_{-i}) \quad \forall a \in A.$$

The pure-strategy Nash equilibria are exactly the fixed points of BR .

By continuity and compactness (=subcompleteness), each BR_i is nonempty-valued, so BR is nonempty-valued. By definition of a supermodular game and Topkis's theorem (p. 60), BR_i is increasing in the induced set order for each $i \in I$; hence BR is increasing in the induced set order. Since each (A_i, \geq) is a subcomplete sublattice of (\mathbf{R}^{m_i}, \geq) , the product lattice (A, \geq) is a subcomplete sublattice of (\mathbf{R}^m, \geq) (where $m := \sum_{i \in I} m_i$), hence a complete lattice in its own right. It follows by Zhou's theorem (p. 59) that the set of fixed points of BR is a nonempty complete lattice. ■




Figure 7 – The lattice (L, \geq) and the increasing function $f : L \times T \rightarrow L$. When t increases, $f(\cdot, t) : L \rightarrow L$ shifts up, causing its greatest and least fixed points to increase.

The other nice thing about supermodular games is their comparative statics properties. The underlying mathematical result is the following.

Proposition 9. If (L, \succsim) is a nonempty complete lattice, (T, \succeq) is a poset, and $f : L \times T \rightarrow L$ is order-preserving, then for each $t \in T$ the set of fixed points of $f(\cdot, t)$ is a nonempty complete lattice with greatest and least elements

$$x^*(t) := \sup_{(L, \succsim)} \{x \in L : x \preceq f(x, t)\} \quad \text{and}$$

$$x_*(t) := \inf_{(L, \succsim)} \{x \in L : x \succeq f(x, t)\},$$

and $x^* : T \rightarrow L$ and $x_* : T \rightarrow L$ are increasing.

Proposition 9 is illustrated in Figure 7.

Proof. Everything except the monotonicity of x_* and x^* is immediate from Tarski's theorem (p. 55). Define

$$R(t) := \{x \in L : x \succeq f(x, t)\},$$

and take $t, t' \in T$ with $t \gtrsim t'$. Since $x_*(t)$ is a fixed point of $f(\cdot, t)$, we have (a fortiori) $x_*(t) \gtrsim f(x_*(t), t)$. Since f is increasing, $f(x_*(t), t) \gtrsim f(x_*(t), t')$. Hence $x_*(t) \gtrsim f(x_*(t), t')$, so $x_*(t) \in R(t')$. Hence

$$x_*(t) \gtrsim \inf_{(L, \gtrsim)} R(t') = x_*(t').$$

The argument for x^* is closely analogous. ■

A parameterised supermodular game is really a collection of supermodular games indexed by the value of a parameter $t \in T$. To obtain comparative statics via Topkis's theorem, we assume that action sets are increasing in t and that payoffs have increasing differences in t .

Definition 22. Let

- (1) I be a set of players,
- (2) (T, \succeq) be a poset parameterising the game,
- (3) \mathcal{A}_i be the set of all actions i might be able to take,
- (4) $A_i : T \Rightarrow \mathcal{A}_i$ return i 's action set, and
- (5) $u_i : A_i \times A_{-i} \times T \rightarrow \mathbf{R}$ be i 's payoff.

$(I, (T, \succeq), \{\mathcal{A}_i, A_i, u_i\}_{i \in I})$ is a parameterised supermodular game iff

- (1) $(I, \{A_i(t), u_i(\cdot, \cdot, t)\}_{i \in I})$ is a supermodular game for each $t \in T$,
- (2) A_i is increasing in the induced order for each $i \in I$, and
- (3) $u_i(\cdot, a_{-i}, \cdot)$ has increasing differences for each $a_{-i} \in \mathcal{A}_{-i}$ and $i \in I$.

The smallest and largest pure-strategy Nash equilibria are increasing:

Corollary 18. Let $(I, (T, \succeq), \{\mathcal{A}_i, A_i, u_i\}_{i \in I})$ be a parameterised supermodular game. For each $t \in T$, let $N(t)$ be the set of pure-strategy Nash equilibria of the supermodular game $(I, \{A_i(t), u_i(\cdot, \cdot, t)\}_{i \in I})$. Then $N(t)$ is a nonempty complete lattice for each $t \in T$, and $t \mapsto \inf_{(N(t), \succeq)} N(t)$ and $t \mapsto \sup_{(N(t), \succeq)} N(t)$ are increasing.

Proof. The fact that $N(t)$ is a nonempty complete lattice for each t is immediate from Proposition 8 (p. 62). The fact that the smallest and largest equilibria are increasing is immediate from Proposition 9 (p. 63). ■

As previously noted, Topkis's theorem remains true if we replace the 'cardinal' notions of supermodularity and increasing differences with their 'ordinal' counterparts, quasi-supermodularity and single crossing. (That's Milgrom and Shannon's (1994) monotonicity theorem). All of the results in this section therefore extend to the class of 'quasi-supermodular games' for which payoffs are quasi-supermodular and single crossing.

Supermodular games have nice properties unrelated to Nash equilibrium. Most importantly, (1) the rationalisable set of a supermodular game lies between the least and greatest pure-strategy Nash equilibria, and (2) the least and greatest Nash equilibria are stable under various evolutionary dynamics. See Milgrom and Roberts (1990).

2.5 Monotone equilibria of supermodular Bayesian games

The ideas in this section are due to van Zandt and Vives (2007). I haven't looked at the paper, so I don't know how similar it is to the presentation below.

Consider the following species of Bayesian game. The players are I , and i 's action space is $A_i \subseteq \mathbf{R}^{m_i}$. i receives a signal taking values in \mathcal{S}_i . Payoffs are given by a measurable map $u_i : A \times \mathcal{S} \rightarrow \mathbf{R}$.³⁹ As in the previous section, assume that for each $i \in I$,

- (1) (A_i, \geq) is a subcomplete sublattice of (\mathbf{R}^{m_i}, \geq) for some $m_i \in \mathbf{N}$,
- (2) $u_i(\cdot, \cdot, s) : A_i \times A_{-i} \rightarrow \mathbf{R}$ has increasing differences for each $s \in \mathcal{S}$, and
- (3) $u_i(\cdot, a_{-i}, s)$ is continuous and supermodular for each $a_{-i} \in A_{-i}$ and $s \in \mathcal{S}$.

Further assume that for each $i \in I$,

- (4) $(\mathcal{S}_i, \succsim)$ is a lattice.
- (5) $u_i(\cdot, a_{-i}, \cdot) : A_i \times \mathcal{S} \rightarrow \mathbf{R}$ has increasing differences for each $a_{-i} \in A_{-i}$.

There's one primitive of a Bayesian game left: i 's belief is a map $\mu_i : \mathcal{S}_i \rightarrow \Delta(\mathcal{S}_{-i})$. Endow $\Delta(\mathcal{S}_{-i})$ with the first-order stochastic dominance (FOSD) order \succsim_F . We know from before that $(\Delta(\mathcal{S}_{-i}), \succsim_F)$ is a lattice. Assume that

- (6) $\mu_i : \mathcal{S}_i \rightarrow \Delta(\mathcal{S}_{-i})$ is increasing.

³⁹It's sometimes more natural to think of payoffs as depending on some state about which the players' signals are informative rather than depending directly on the signals. This setup can be derived from such a setting.

We can now establish the existence of monotone pure-strategy Bayes–Nash equilibria. Suppose that players $-i$ are playing an increasing pure strategy profile $\sigma_{-i} : \mathcal{S}_{-i} \rightarrow A_{-i}$. If i takes action a_i and signals are (s_i, s_{-i}) then her payoff is

$$U_i^{\sigma_{-i}}(a_i, s_{-i}, s_i) := u_i(a_i, \sigma_{-i}(s_{-i}), s_i, s_{-i}).$$

So her expected payoff from playing a_i after observing signal s_i is

$$\mathcal{U}_i^{\sigma_{-i}}(a_i, s_i) := \int_{\mathcal{S}_{-i}} U_i^{\sigma_{-i}}(a_i, s_{-i}, s_i) \mu_i(s_i) (ds_{-i}).$$

$\mathcal{U}_i^{\sigma_{-i}}(\cdot, s_i)$ is supermodular for each $s_i \in \mathcal{S}_i$ since $u_i(\cdot, a_{-i}, s_i, s_{-i})$ is. The existence of an increasing best response therefore hinges on whether $\mathcal{U}_i^{\sigma_{-i}}$ has increasing differences. The following (trivial) lemma tells us that it does:

Lemma 6. Let (L, \succsim) , (L', \succsim') and (L'', \succsim'') be lattices, and let $\phi : L \times L' \times L'' \rightarrow \mathbf{R}$ and $\nu : L'' \rightarrow \Delta(L')$.⁴⁰ Assume that $\phi(\cdot, \cdot, z)$ for each $z \in L''$ and $\phi(\cdot, y, \cdot)$ for each $y \in L'$ have increasing differences and that ν is increasing in the FOSD order. Then $\Psi : L \times L'' \rightarrow \mathbf{R}$ defined by

$$\Phi(x, z) := \int_{L'} \phi(x, y, z) \nu(z) (dy)$$

has increasing differences.

Proof. Take $x \geq x'$ in L and $z \geq z'$ in L'' . We have

$$\begin{aligned} \Phi(x, z) - \Phi(x', z) &= \int_{L'} [\phi(x, y, z) - \phi(x', y, z)] \nu(z) (dy) \\ &\geq \int_{L'} [\phi(x, y, z') - \phi(x', y, z')] \nu(z) (dy) \\ &\geq \int_{L'} [\phi(x, y, z') - \phi(x', y, z')] \nu(z') (dy) \\ &= \Phi(x, z') - \Phi(x', z'), \end{aligned}$$

where the first inequality holds since $\phi(\cdot, y, \cdot)$ has increasing differences, and the second inequality used the fact that $\phi(\cdot, \cdot, z)$ has increasing differences and that $\nu(z)$ first-order stochastically dominates $\nu(z')$. \blacksquare

$U_i^{\sigma_{-i}}(\cdot, \cdot, s_i)$ has increasing differences since u_i does and σ_{-i} is monotone. $U_i^{\sigma_{-i}}(\cdot, s_{-i}, \cdot)$ has increasing differences since u_i does. μ_i is increasing in FOSD. Hence by Lemma 6, $\mathcal{U}_i^{\sigma_{-i}}$ has increasing differences.

⁴⁰There's a σ -algebra in here as well, of course.

Since $\mathcal{U}_i^{\sigma_{-i}}$ has increasing differences, i 's the best response to a strategy σ_{-i} is increasing in her signal by Topkis's monotonicity theorem (p. 60). So there is a monotone best reply σ_i to any monotone strategy profile σ_{-i} .

We can now establish the existence of monotone Bayes–Nash equilibria using Tarski's fixed-point theorem again. Simply observe that since there are monotone best responses to monotone profiles, the Nash correspondence restricted to monotone strategies is well-defined, nonempty-valued and increasing, hence has a fixed point by Tarski's fixed-point theorem.

Since we used Tarski's fixed-point theorem to prove existence, we can also recover the nice properties of complete-information supermodular games from the previous section. In particular, the set of monotone equilibria has a lattice structure, and if we parametrise the game appropriately then the largest and smallest equilibria will be increasing in the parameter.

What if we want to weaken supermodularity and increasing differences to quasi-supermodularity and single crossing? The only obstacle is that we need a result along the lines of Lemma 6 to ensure that single crossing is preserved by integration. There is such a result, but it requires a stronger order on the distributions.

So endow $\Delta(\mathcal{S}_{-i})$ with the monotone likelihood ratio property (MLRP) order \succsim_M , defined as follows. Fix a reference measure λ with respect to which $\{\{\mu_i(s_i)\}_{s_i \in \mathcal{S}_i}\}_{i \in I}$ are absolutely continuous; then $\mu \succsim_M \mu'$ iff the likelihood ratio

$$\frac{d\mu/d\lambda}{d\mu'/d\lambda}$$

is increasing. $(\Delta(\mathcal{S}_{-i}), \succsim_M)$ is a lattice. Assume as before that

$$(6) \quad \mu_i : \mathcal{S}_i \rightarrow \Delta(\mathcal{S}_{-i}) \text{ is increasing.}$$

The assumption looks the same, but ‘increasing’ means something else now. It is in fact a stronger assumption: it's easily verified that $\mu \succsim_M \mu'$ implies $\mu \succsim_F \mu'$.

The MLRP order is sufficiently strong that if we weaken ‘increasing differences’ to ‘single crossing’ in Lemma 6, but strengthen ‘FOSD’ to ‘MLRP’, then the result continues to hold. This (nontrivial) result is due to Karlin and Rubin (1956).⁴¹ It follows that ‘quasi-supermodular Bayesian games’ whose payoffs are quasi-supermodular and single crossing and whose interim

⁴¹The MLRP order (and many of its properties) are due to the statistician Samuel Karlin. Most of the results useful for economics are from Karlin and Rubin (1956); Karlin and Rinott (1980) is a nice survey with cleaner proofs. MLRP is usually called ‘total positivity’ in statistics. The MLRP was first introduced into economics by Milgrom (1981).

beliefs are MLRP-increasing have the same properties derived above: the set of monotone Bayes–Nash equilibria is nonempty with a lattice structure, and is increasing if we parameterise appropriately.

Observe that μ being MLRP-increasing is equivalent to $\mu(\cdot)(\cdot)$ being log-supermodular (i.e. $\ln \mu(\cdot)(\cdot)$ being supermodular). It's easy to show that if the joint density of all players' signals is log-supermodular, then the conditional densities are log-supermodular, hence MLRP-increasing. Random variables whose joint density is log-supermodular are called affiliated. Affiliation plays an important role in the analysis of general auction games because it (together with quasi-supermodularity and single crossing of payoffs) ensures equilibrium existence. (And other nice properties, as we saw.) Milgrom and Weber (1982) is the seminal paper, and Krishna (2010, ch. 6) provides a textbook treatment.

An extension to dynamic Bayesian games can be found in Mensch (2016), who gives conditions for the existence of monotone perfect Bayesian equilibria.

3 Learning

A variety of literatures in game theory are concerned with learning in strategic settings.

One is repeated games with one-sided incomplete information (Aumann & Maschler, 1995). Here the uninformed player can learn about the state by observing the informed player's actions, and the informed player must take this into account when deciding how to make use of her private information.

Another is reputation (Fudenberg & Levine, 1989, 1992). There is a long-run player who plays a stage game against a sequence of short-lived players. The long-run player may be of a 'crazy' type which always plays a certain strategy. The short-lived players observe the history of play, and use this to learn about the long-run player's type. The long-run player therefore has an incentive to play so as to appear to be crazy even when she is not.⁴² Such a model can, for example, be used to explain why real-life bargaining often involves delay (Abreu & Gül, 2000).

Yet another is learning to play Nash equilibrium in Bayesian games (Kalai & Lehrer, 1993). Players repeatedly play a Bayesian game. By observing each others' actions, they learn about each others' types. We might then hope that play will converge to what it would be if the types were common knowledge, viz. a Nash equilibrium of the complete-information game. Kalai and Lehrer (1993) provide conditions under which this occurs.

3.1 Repeated games with one-sided incomplete information

In a repeated game with incomplete information on one side, one player knows the payoffs and the other does not. However, the uninformed player can learn about the payoffs by observing the informed player's actions. The informed player therefore has to take into account that her actions will reveal some of her private information to her opponent. What payoff can the informed player achieve in such a setting, and what strategy does he use to achieve it?

This question was answered for zero-sum games by Aumann and Maschler (1995).⁴³ A textbook treatment can be found in Sorin (2002, ch. 3).

⁴²In my opinion, 'reputation' is not a very good name for this phenomenon.

⁴³The paper was first circulated in 1966.

3.1.1 Environment

A two-player zero-sum game (in normal form) is defined as follows. There are two players, Xaver and Yulia. Xaver takes an action $x \in \mathcal{X}$, and Yulia takes an action $y \in \mathcal{Y}$. There is a function $u : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbf{R}$, and Xaver's payoff is u while Yulia's is $-u$. We consider only finite games, meaning that \mathcal{X} and \mathcal{Y} are finite.

Strategies are distributions $\sigma \in \Delta(\mathcal{X})$ and $\tau \in \Delta(\mathcal{Y})$. The value of the game $(\mathcal{X}, \mathcal{Y}, u)$ is defined

$$v := \min_{\tau \in \Delta(\mathcal{Y})} \max_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} u(x, y) \tau(y) = \max_{\sigma \in \Delta(\mathcal{X})} \min_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} u(x, y) \sigma(x),$$

where the equality holds by a minmax theorem. Evidently any Nash equilibrium must involve minmax strategies, and must yield payoffs $(v, -v)$.

Now suppose a finite two-player zero-sum game is played repeatedly, and that Xaver knows the payoffs u but that Yulia does not. In particular, let \mathcal{S} be a finite set of states of nature. Xaver knows the state, whereas Yulia has a belief $p \in \Delta(\mathcal{S})$ about it. Payoffs depend on the state: u maps $\mathcal{X} \times \mathcal{Y} \times \mathcal{S}$ into \mathbf{R} . The game is played $n \in \mathbf{N}$ times, and payoffs are the average of stage game payoffs. We're interested in play and payoffs for n large.

The history of play prior to period t is an element of $(\mathcal{X} \times \mathcal{Y})^{t-1}$, where we use the convention $(\mathcal{X} \times \mathcal{Y})^0 := \{\emptyset\}$. Behavioural strategies of Xaver and Yulia are, respectively, maps

$$\sigma : \mathcal{S} \times \bigcup_{t=0}^{n-1} (\mathcal{X} \times \mathcal{Y})^t \rightarrow \Delta(\mathcal{X}) \quad \text{and} \quad \tau : \bigcup_{t=0}^{n-1} (\mathcal{X} \times \mathcal{Y})^t \rightarrow \Delta(\mathcal{Y}).$$

By Kuhn's (1953) theorem, it is without loss to restrict attention to behavioural strategies.

If Yulia's belief were fixed at p , then it would be easy to work out how Xaver should make optimal use of his information—the problem is essentially nonstrategic. But in equilibrium, Yulia understands how Xaver behaves in different states, and so she learns about the state by observing Xaver's actions. As a result, it is nontrivial to figure out how Xaver should use his information optimally and what value he is thus able to obtain.

3.1.2 The theorem

To get a feeling for what Xaver can achieve by using his private information in different ways, let's look at a few examples.

Example 11. Consider $\mathcal{X} = \{T, B\}$, $\mathcal{Y} = \{L, R\}$, $\mathcal{S} = \{1, 2\}$ and $p(1) = p(2) = 1/2$, and let u be given by

$s = 1$		$s = 2$			
	L	R	L	R	
T	1	0	T	0	0
B	0	0	B	0	1

Suppose that Xaver uses a pure strategy. Then Yulia learns the state right away, and can subsequently ensure that the payoff is 0 forevermore. This isn't very good for Xaver!

Suppose instead that Xaver mixes 50–50 at every history, ignoring his private information. Then the game is essentially a static zero-sum game with payoffs

	L	R
T	1/2	0
B	0	1/2

The value of this game is 1/2. So in this example, Xaver does better by ignoring his private information. (It can be shown that this is actually Xaver's optimal strategy in this game.)

Example 12. Consider $\mathcal{X} = \{T, B\}$, $\mathcal{Y} = \{L, R\}$, $\mathcal{S} = \{1, 2\}$ and $p(1) = p(2) = 1/2$, but let payoffs u be

$s = 1$		$s = 2$			
	L	R	L	R	
T	-1	0	T	0	0
B	0	0	B	0	-1

If Xaver plays B whenever the state is 1 and T whenever the state is 2 (at every history), he ensures a payoff of 0. He obviously cannot do better. So in this example, Xaver finds it optimal to reveal his information.

Example 13. There are intermediate cases in which Xaver finds it optimal 'partially reveal' his private information. Consider $\mathcal{X} = \{T, B\}$, $\mathcal{Y} = \{L, R, M\}$, $\mathcal{S} = \{1, 2\}$ and $p(1) = p(2) = 1/2$, with payoffs

$s = 1$		$s = 2$					
	L	M	R	L	M	R	
T	4	0	2	T	0	4	-2
B	4	0	-2	B	0	4	2

If Xaver uses a pure strategy and thereby reveals his private information, Yulia will subsequently play L forever if the state is 1 and M forever if the state is 2. Either way, the payoff is subsequently zero forever.

If Xaver reveals nothing by mixing 50–50 regardless of his information, the game effectively becomes the static game

	L	M	R
T	2	2	0
B	2	2	0

Yulia will play R , so the payoff will again be 0.

Consider instead the following strategy for Xaver. In state 1, w.p. $3/4$ play T forever, and w.p. $1/4$ play B forever; in state 2, w.p. $1/4$ play T forever, and w.p. $3/4$ play B forever. Yulia observes what Xaver plays and updates her belief about the state accordingly. Her conditional expected payoffs are

	T		B		
L	M	R	L	M	R
3	1	1	1	3	1

So R is always optimal for Yulia. But if Yulia always plays R , then this strategy yields ex ante expected payoff $1 > 0$ for Xaver. (In fact, this is Xaver’s optimal strategy in this example.)

Now let’s be a little more formal. Let $v_n(p)$ be the value of the repeated game when Yulia’s belief is $p \in \Delta(\mathcal{S})$. We can think of v_n as what Xaver can achieve by making optimal use of his private information. In the three examples above, we studied v_n as a function of the fundamentals.

Let $f(p)$ be the value of the static game in which neither player knows the state, and their belief about it is $p \in \Delta(\mathcal{S})$; this is what Xaver can achieve by ignoring his private information. Let $\text{cav } f$ be the concavification of f , i.e. the smallest concave function weakly above f .⁴⁴ We have the following beautiful characterisation of v_n :

Theorem 20 (Aumann and Maschler (1995)). There is a constant $C > 0$ such that $\text{cav } f \leq v_n \leq \text{cav } f + n^{-1/2}C$ for all $n \in \mathbf{N}$.

It follows in particular that $v_n(p) \searrow \text{cav } f(p)$ for every $p \in \Delta(\mathcal{S})$.

To understand what this means, suppose that $\mathcal{S} = \{1, 2\}$ as in the examples above. We can then think of $p \in \Delta(\mathcal{S})$ as a single number $p(1) =: p$ in $[0, 1]$, and hence define f on $[0, 1]$. An example of such an f and its concavification are drawn in Figure 8.

⁴⁴C.f. the convexification operator vex from section 1.2.8 (p. 25).




Figure 8 – A function $f : [0, 1] \rightarrow \mathbf{R}$ (solid) and its concavification (dotted).

The intuition for the first inequality ($\text{cav } f \leq v_n$) in the theorem is as follows. Xaver can always achieve f by ignoring his information (as was optimal in Example 11). He can also achieve the convex combination $(1-p)f(0) + pf(1)$ by playing a pure strategy (as was optimal in Example 12). By using state-dependent randomisations of the form we saw in Example 13, he can achieve other convex combinations. In Figure 8, Xaver can achieve the payoffs on the dotted line segment by using state-dependent mixtures that induce beliefs $1/4$ and $7/8$ in Yulia with some probabilities.

The second inequality ($v_n \leq \text{cav } f + n^{-1/2}C$) is less intuitive, and harder to prove. The basic idea is that the manipulation above (using convex combinations of beliefs) is in fact the only kind of manipulation that Xaver can profitably engage in in the long run.

It is straightforward to verify that the long-run values $\lim_{n \rightarrow \infty} v_n$ in Examples 11 to 13 are equal to $\text{cav } f(1/2)$ in each case. In Example 11, where Xaver found it optimal to ignore his information, f is concave, so the value is the uninformed value $f(1/2)$, and Xaver performs no state-dependent mixing. In Example 12, where Xaver found it optimal to reveal the state, f is convex, so v_n is always above f , and he finds it optimal to induce the beliefs $p = 0$

and $p = 1$ (depending on the state). In Example 13, f is neither concave nor convex, and for prior $1/2$ it is optimal for Xaver to induce beliefs $1/4$ and $3/4$.

3.1.3 The lower bound on v_n

The intuition we provided for why v_n must be at least $\text{cav } f$ was that Xaver can randomly induce posteriors for Yulia to guarantee $\text{cav } f$. The following lemma provides the formal sense in which that is the case.

Lemma 7 (splitting lemma). For any finite set \mathcal{L} , $\{\lambda_\ell\}_{\ell \in \mathcal{L}} \subseteq [0, 1]$ and $\{p_\ell\}_{\ell \in \mathcal{L}} \subseteq \Delta(\mathcal{S})$ which satisfy

$$\sum_{\ell \in \mathcal{L}} p_\ell \lambda_\ell = p,$$

there exists a random variable L taking values in \mathcal{L} such that the joint distribution \mathbf{P} of the signal L and the state S satisfies $\mathbf{P}(L = \ell) = \lambda_\ell$ and $\mathbf{P}(S = s | L = \ell) = p_\ell(s)$ for every $(\ell, s) \in \mathcal{L} \times \mathcal{S}$.

The converse is obvious: for any signal S jointly distributed with L , the martingale property $\sum_{\ell \in \mathcal{L}} p_\ell \lambda_\ell = p$ of posteriors must hold. The splitting lemma is usually attributed to Aumann and Maschler (1995), but it can be traced back to Blackwell (1951).

Proof. Let L be the \mathcal{L} -valued random variable with conditional distribution

$$\mathbf{P}(L = \ell | S = s) := \frac{p_\ell(s) \lambda_\ell}{p(s)} \quad \forall (\ell, s) \in \mathcal{L} \times \mathcal{S}.$$

For any $\ell \in \mathcal{L}$, by the law of total probability,

$$\mathbf{P}(L = \ell) = \sum_{s \in \mathcal{S}} \mathbf{P}(L = \ell | S = s) \mathbf{P}(S = s) = \sum_{s \in \mathcal{S}} \frac{p_\ell(s) \lambda_\ell}{p(s)} p(s) = \lambda_\ell.$$

For any $(\ell, s) \in \mathcal{L} \times \mathcal{S}$, by Bayes's rule,

$$\mathbf{P}(S = s | L = \ell) = \frac{\frac{p_\ell(s) \lambda_\ell}{p(s)} p(s)}{\lambda_\ell} = p_\ell(s). \quad \blacksquare$$

To link the splitting lemma to concavification, we use the following.

Lemma 8. Consider $f : \Delta(\mathcal{S}) \rightarrow \mathbf{R}$.

(1) For every $p \in \Delta(\mathcal{S})$,

$$\text{cav } f(p) = \sup_{\mathcal{L}, \{\lambda_\ell, p_\ell\}_{\ell \in \mathcal{L}}} \left\{ \sum_{\ell \in \mathcal{L}} f(p_\ell) \lambda_\ell : \sum_{\ell \in \mathcal{L}} p_\ell \lambda_\ell = p \right\}.$$

(2) The equation continues to hold if we require \mathcal{L} to have at most $|\mathcal{S}| + 1$ elements: for every $p \in \Delta(\mathcal{S})$,

$$\text{cav } f(p) = \sup_{\{\lambda_\ell, p_\ell\}_{\ell=1}^{|\mathcal{S}|+1}} \left\{ \sum_{\ell=1}^{|\mathcal{S}|+1} f(p_\ell) \lambda_\ell : \sum_{\ell=1}^{|\mathcal{S}|+1} p_\ell \lambda_\ell = p \right\}.$$

(3) If f is upper semicontinuous, then the supremum is attained.

Properties (1) and (2) are essentially equivalent to Carathéodory's theorem. Property (3) is immediate from Berge's theorem.⁴⁵

With these properties in hand, $\text{cav } f \leq v_n$ is (almost) immediate. By the splitting lemma, Xaver can attain any value $\sum_\ell f(p_\ell) \lambda_\ell$ for which the distribution $\{\lambda_\ell\}$ over posteriors $\{p_\ell\}$ of Yulia satisfies the martingale property. (He does this by splitting himself in period zero.) By Lemma 8, his value must therefore be at least $\text{cav } f(p)$.

This splitting argument has seen recent use in different setting:

Example 14 (persuasion). An informed principal (Xaver) knows the state in \mathcal{S} , whereas the agent (Yulia) has only a belief $p \in \Delta(\mathcal{S})$. The principal sends a signal to influence the agent's belief and hence her action. The agent observes the signal, then takes the action. The action is noncontractible, and there are no transfers.

An information policy of the principal is a conditional distribution of a signal that varies with the state. Following Kamenica and Gentzkow (2011), assume that the principal can commit to an information policy. This assumption is what distinguishes the persuasion literature from the older literature on strategic communication, notably cheap talk (Crawford & Sobel, 1982) and verifiable disclosure (Grossman, 1981; Milgrom, 1981).

Let $f(p)$ be the payoff to the principal from the action that the agent takes when her belief is p . Assume upper semicontinuity of f —this means precisely that the agent breaks ties in the principal's favour.

⁴⁵We mentioned these properties when discussing the convexification operator vex in section 1.2.8 (p. 25).




Figure 9 – f (solid) and its concavification (dotted) in the trial example.

By the splitting lemma, the principal can induce a distribution $\{\lambda_\ell\}$ over posteriors $\{p_\ell\}$ iff they satisfy the martingale property $\sum_\ell p_\ell \lambda_\ell = p$ of posteriors. Hence the expected payoffs that the principal can achieve are precisely the $\sum_\ell f(p_\ell) \lambda_\ell$ which satisfy $\sum_\ell p_\ell \lambda_\ell = p$. By Lemma 8, it follows that the principal’s maximal value is equal to $\text{cav } f(p)$.

There is a difference here from the repeated games setting, where the splitting lemma only gave us that the value is $\geq \text{cav } f(p)$ (not necessarily equal). The reason is that in a repeated game, Xaver’s strategies differ in respects other than what distribution over posteriors they induce in Yulia.

To illustrate, consider the example of a trial. The state is whether a defendant is guilty. A judge (agent, Yulia) believes with probability p that the defendant is guilty, and must convict or not convict her. The judge gets 1 iff she makes the just decision, so she convicts iff her belief is $\geq 1/2$. The DA (principal, Xaver) carries out an investigation and reports all of its findings to the judge. (This is the commitment assumption.) He gets 1 iff the defendant is convicted, so designs his investigation to maximise the probability of conviction.

The DA’s payoff f is given by $f(p) = \mathbf{1}(p \geq 1/2)$, with concavification $\text{cav } f(p) = \min\{2p, 1\}$. f and $\text{cav } f$ are depicted in Figure 9. Focus on the interesting case $p < 1/2$. From the figure, it’s clear that the optimal signal always induces the lowest posterior belief consistent with the agent taking a given action, in this case 0 and $1/2$. This is optimal because inducing favourable (high) beliefs is costly: the martingale property $\sum_\ell p_\ell \lambda_\ell = p$ acts like a budget constraint, so that inducing a higher belief p_ℓ with some probability will require a lower belief $p_{\ell'}$ with some probability. See Kamenica and Gentzkow (2011) for a precise statement of how this property generalises

beyond the trial example.

Finding the optimal information policy is straightforward in the trial example. Feasible distributions $(1 - \lambda, \lambda)$ over posteriors 0 and $1/2$ are those that satisfy the martingale property $0 \cdot (1 - \lambda) + \frac{1}{2}\lambda = p$, or $\lambda = 2p$. Let the signal value that induces belief 0 (hence no conviction) be called i , and the value that induces $1/2$ (hence conviction) be called g . By the splitting lemma, the distribution of the signal L given the state S must be

$$1 - \mathbf{P}(L = i | S = \text{innocent}) = \mathbf{P}(L = g | S = \text{innocent}) = \frac{\frac{1}{2} \cdot \lambda}{1 - p} = \frac{p}{1 - p}$$

$$1 - \mathbf{P}(L = i | S = \text{guilty}) = \mathbf{P}(L = g | S = \text{guilty}) = \frac{\frac{1}{2} \cdot \lambda}{p} = 1.$$

So the optimal investigation is such that if the defendant is guilty then the investigation discovers this with probability 1, whereas if the defendant is innocent then the investigation only finds this out with probability $p/(1 - p)$.

When this information policy is used, the judge convicts the defendant with probability $\mathbf{P}(L = g) = \frac{p}{1-p} \cdot (1 - p) + 1 \cdot p = 2p$. This is despite the fact that the judge only believes the defendant to be guilty with probability $p < 2p$! This is perhaps counterintuitive: it looks like the agent is being systematically fooled. But she is not—indeed, since the principal commits to his information policy, there is no way for the principal to fool the agent either on or off the equilibrium path.

3.1.4 The upper bound on v_n

We showed that v_n must be at least $\text{cav } f$. There wasn't much to that result: once you understand the splitting lemma, it's almost immediate.

Showing that v_n cannot be much larger than $\text{cav } f$ is harder. The idea of the proof is that after the initial randomisation that is used to concavify f , Yulia's beliefs cannot be induced to vary that much, so Xaver's payoff cannot improve much. In particular, for any strategy of Xaver, Yulia has a strategy that forces the payoff to be at most $\text{cav } f + n^{-1/2}C$.

To formalise the notion that Yulia's beliefs can only change so much, we first need to decide what we mean by two beliefs being 'close'. We're going to use the ℓ^1 norm: the distance between $q \in \Delta(\mathcal{S})$ and $q' \in \Delta(\mathcal{S})$ is

$$\|q - q'\|_{\ell^1} := \sum_{s \in \mathcal{S}} |q(s) - q'(s)|.$$

Some equivalent ways of expressing this metric are

$$\|q - q'\|_{\ell^1} = \sup_{g \in [0,1]^{\mathcal{S}}} \left| \int_{\mathcal{S}} g d(q - q') \right| = 2 \sup_{\substack{B \subseteq \mathcal{S} \\ \text{measurable}}} |q(B) - q'(B)|.$$

(This isn't hard to prove.)

We use the ℓ^1 norm because it is payoff-relevant. In particular, consider a decision taken under uncertainty, where the uncertainty is actually distributed according to $q \in \Delta(\mathcal{S})$. Write $V(q')$ for the value obtained by a decision-maker who takes the decision that is optimal under $q' \in \Delta(\mathcal{S})$. Then if $\|q - q'\|_{\ell^1} < \varepsilon$, we have $|V(q) - V(q')| < 2\varepsilon$.

Now consider a bounded martingale $\{q_t\}$ on \mathbf{R}^m , where q_1 is constant a.s. For any norm $\|\cdot\|$ on \mathbf{R}^m , the total $\|\cdot\|$ -variation of $\{q_t\}$ is defined

$$V_n^{\|\cdot\|}(\{q_t\}) := \mathbf{E} \left(\sum_{t=1}^n \|q_{t+1} - q_t\| \right).$$

We know that martingales converge, so that there is a bound on the rate of growth of the total variation in n . The following tells us that for the ℓ^1 norm, that bound is $O(n^{1/2})$.⁴⁶

Lemma 9 (Mertens and Zamir (1977)). For any a bounded martingale $\{q_t\}$ on \mathbf{R}^m , there is a constant $M > 0$ s.t. $V_n^{\ell^1}(\{q_t\}) \leq n^{1/2}M$ for every $n \in \mathbf{N}$.

A nice proof can be found in Neyman (2013).

We can now provide a rough idea of why Xavier cannot do much better than $\text{cav } f$. Since the total variation of Yulia's belief martingale is bounded by $O(n^{1/2})$, in most periods, her belief doesn't change much. When Yulia's belief is nearly unchanged, the payoff is nearly what it would be if her belief were fixed forever after the initial period (like in the persuasion problem). That payoff is $\text{cav } f(p)$, as we know.

A slightly more formal argument is as follows. Xavier begins by splitting himself, so that the value would be $\text{cav } f(p)$ if Yulia stopped learning. The Mertens–Zamir lemma tells us that Yulia can only learn so much: her belief changes by $> \varepsilon$ in at most $Mn^{1/2}/\varepsilon$ of the first n periods. Let $k(\varepsilon)$ be the largest payoff increase achievable in one period. Then

$$\begin{aligned} v_n &\leq \text{cav } f(p) + n^{-1} \left[\frac{Mn^{1/2}}{\varepsilon} \cdot k(\varepsilon) + \left(1 - \frac{Mn^{1/2}}{\varepsilon} \right) \cdot 2\varepsilon \right] \\ &= \text{cav } f(p) + n^{-1/2}M \left(\frac{k(\varepsilon)}{\varepsilon} - 2 \right) + O(\varepsilon). \end{aligned}$$

⁴⁶Similar results are available for other norms, e.g. the Kullback–Leibler norm.

The trick is to show that $k(\varepsilon)/\varepsilon$ converges to a constant as $\varepsilon \rightarrow 0$.

The broad message is the use of the Mertens–Zamir lemma to bound payoffs. The idea is that in the long run, in most periods, Yulia can’t learn much more. Play will therefore mostly be as if neither player knew the state. (In our case, payoffs under two-sided incomplete information was *cav f*.) Even though Yulia need never learn the state (and Xaver certainly does not forget it), play is roughly as though information were symmetric.

This is not the argument originally used by Aumann and Maschler (1995). (The Mertens–Zamir lemma was published a decade later, after all.) They instead show only the limit result that v_n must converge to *cav f*. The real accomplishment is that they construct Yulia’s optimal strategy, in sharp contrast to the nonconstructive argument that we gave. Their construction is based on Blackwell’s (1956) approachability theorem.

3.2 Reputation and learning to play Bayesian games

A result akin to the Mertens–Zamir lemma is also the driving force in papers on reputation in games. This literature started with Kreps, Milgrom, Roberts and Wilson (1982) and Fudenberg and Levine (1989, 1992).

To illustrate, consider the stage game

	<i>L</i>	<i>R</i>
<i>T</i>	0, 0	4, 4
<i>B</i>	0, 0	5, -10

Suppose that the columns are played by a sequence of short-lived players who observe the history of play, whereas the rows are played by an ordinary long-run player. There is payoff uncertainty: with some (small) probability, the long-run player is of a ‘crazy’ type that always plays *T*. The non-crazy type of the long-run player has an incentive to pool with the crazy type (to ‘maintain a crazy reputation’), since this will induce the short-lived players to choose *R*.

A Mertens–Zamir-type result implies that the beliefs of the short-run players cannot change that much in most periods. In particular, if the state is S and actions are (A_1, A_2, \dots) , the result says that

$$\mathbf{P}(S|A_1, \dots, A_n) \quad \text{and} \quad \mathbf{P}(S|A_1, \dots, A_{n+1})$$

can’t be very far apart. The metric here is the Kullback–Leibler metric. Because this metric is symmetric in S and A_{n+1} , the Mertens–Zamir-type result is equivalent to the result that

$$\mathbf{P}(A_{n+1}|A_1, \dots, A_n, S) \quad \text{and} \quad \mathbf{P}(A_{n+1}|A_1, \dots, A_n)$$

can't be very far apart. This says in most periods, the prediction of the short-run player wouldn't be very different even if she knew the type of the long-run player! The short-run players need not learn the type of the long-run player: it is just that they can form very good predictions of the long-run player's future behaviour without knowing the state.

This reasoning allows us to bound the long-run player's payoff across all equilibria. This is an unusually strong result: usually in repeated games, there are almost no bounds on what payoffs can be supported in equilibrium! (That's what folk theorems say.)

Also related is the literature on (Bayesian) learning in games Kalai and Lehrer (1993). Players repeatedly play a game of incomplete information, and learn about each others' private information by observing each others' actions. By a Mertens–Zamir-type result, although players need never learn each others' information, they will eventually learn to make very accurate predictions about each others' play. It follows that play will converge (in a particular sense, under some conditions) to a Bayes–Nash equilibrium of the stage game.

Sorin (1999) provides two general Mertens–Zamir-type theorems that unify the three literatures above: repeated games of incomplete information, reputation, and learning in games.⁴⁷

⁴⁷There are other literatures on learning in strategic settings which fall outside the scope of Sorin's theorems. An obvious example is the strategic experimentation literature; see Hörner and Skrzypacz (2016) for a survey.

References

- Abreu, D. & Gül, F. (2000). Bargaining and reputation. *Econometrica*, 68(1), 85–117.
- Abreu, D., Pearce, D., & Stacchetti, E. (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5), 1041–1063.
- Ashlagi, I., Braverman, M., Hassidim, A., & Monderer, D. (2010). Monotonicity and implementability. *Econometrica*, 78(5), 1749–1772.
- Aumann, R. J. & Maschler, M. B. (1995). *Repeated games with incomplete information*. Cambridge, MA: MIT Press.
- Bergemann, D. & Morris, S. (2005). Robust mechanism design. *Econometrica*, 73(6), 1771–1813.
- Bikhchandani, S., Chatterji, S., Lavi, R., Mu’alem, A., Nisan, N., & Sen, A. (2006). Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4), 1109–1132.
- Blackwell, D. (1951). Comparison of experiments. In J. Neyman (Ed.), *Proceedings of the second Berkeley symposium on mathematical statistics and probability* (pp. 93–102). Berkeley, CA: University of California Press.
- Blackwell, D. (1956). An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1), 1–8.
- Border, K. C. (1991). Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59(4), 1175–1187.
- Brouwer, L. E. J. (1912). Über Abbildung von Mannigfaltigkeiten. *Mathematische Annalen*, 71(1), 97–115.
- Bulow, J. & Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1), 180–194.
- Chung, K.-S. & Ely, J. C. (2006). *Ex-post incentive compatible mechanism design*. Working paper, 17 May 2006.
- Crawford, V. P. & Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6), 1431–1451.
- Crémer, J. & McLean, R. P. (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6), 1247–1257.
- Diamond, P. A. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy*, 90(5), 881–894.
- Fudenberg, D. & Levine, D. K. (1989). Reputation and equilibrium selection in games with a patient player. *Econometrica*, 57(4), 759–778.

- Fudenberg, D. & Levine, D. K. (1992). Maintaining a reputation when strategies are imperfectly observed. *Review of Economic Studies*, 59(3), 561–579.
- Gale, D. & Shapley, L. S. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1), 9–15.
- Gershkov, A., Goeree, J. K., Kushnir, A., Moldovanu, B., & Shi, X. (2013). On the equivalence of Bayesian and dominant strategy implementation. *Econometrica*, 81(1), 197–220.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, 24(3), 461–483.
- Gutmann, S., Kemperman, J. H. B., Reeds, J. A., & Shepp, L. A. (1991). Existence of probability measures with given marginals. *Annals of Probability*, 19(4), 1781–1797.
- Hart, S. & Reny, P. J. (2015). Maximal revenue with multiple goods: Non-monotonicity and other observations. *Theoretical Economics*, 10(3), 893–922.
- Heifetz, A. & Neeman, Z. (2006). On the generic (im)possibility of full surplus extraction in mechanism design. *Econometrica*, 74(1), 213–233.
- Hörner, J. & Skrzypacz, A. (2016). *Learning, experimentation and information design*. Working paper, 4 Aug 2016.
- Kalai, E. & Lehrer, E. (1993). Rational learning leads to Nash equilibrium. *Econometrica*, 61(5), 1019–1045.
- Kamenica, E. & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590–2615.
- Karlin, S. & Rinott, Y. (1980). Classes of orderings of measures and related correlation inequalities. I. Multivariate totally positive distributions. *Journal of Multivariate Analysis*, 10(4), 467–498.
- Karlin, S. & Rubin, H. (1956). The theory of decision procedures for distributions with monotone likelihood ratio. *Annals of Mathematical Statistics*, 27(2), 272–299.
- Knuth, D. E. (1976). *Mariages stables et leurs relations avec d'autres problèmes combinatoires : Introduction à l'analyse mathématique des algorithmes*. Montréal : Les Presses de l'Université de Montréal.
- Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, 27(2), 245–252.
- Krishna, V. (2010). *Auction theory* (2nd). Amsterdam: Academic Press.
- Kuhn, H. W. (1953). Extensive games and the problem of information. In H. W. Kuhn & A. W. Tucker (Eds.), *Contributions to the theory*

- of games (Chap. 11, Vol. 2, pp. 193–216). Princeton, NJ: Princeton University Press.
- Manelli, A. M. & Vincent, D. R. (2010). Bayesian and dominant-strategy implementation in the independent private-values model. *Econometrica*, 78(6), 1905–1938.
- McAfee, R. P. & Reny, P. J. (1992). Correlated information and mechanism design. *Econometrica*, 60(2), 395–421.
- Mensch, J. (2016). *On the existence of monotone pure-strategy perfect Bayesian equilibrium in games with complementarities*. Working paper.
- Mertens, J.-F. & Zamir, S. (1977). The maximal variation of a bounded martingale. *Israel Journal of Mathematics*, 27(3–4), 252–276.
- Milgrom, P. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2), 380–391.
- Milgrom, P. & Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complements. *Econometrica*, 58(6), 1255–1277.
- Milgrom, P. & Shannon, C. (1994). Monotone comparative statics. *Econometrica*, 62(1), 157–180.
- Milgrom, P. & Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50(5), 1089–1122.
- Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38(2), 175–208.
- Mirrlees, J. A. (1976). Optimal tax theory: A synthesis. *Journal of Public Economics*, 6(4), 327–358.
- Mussa, M. & Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18(2), 301–317.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1), 58–73.
- Myerson, R. B. & Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2), 265–281.
- Neeman, Z. (2004). The relevance of private information in mechanism design. *Journal of Economic Theory*, 117(1), 55–77.
- Neyman, A. (2013). The maximal variation of martingales of probabilities and repeated games with incomplete information. *Journal of Theoretical Probability*, 26(2), 557–567.
- Roberts, K. (1979). The characterization of implementable choice rules. In J.-J. Laffont (Ed.), *Aggregation and revelation of preferences* (Vol. 6A, pp. 312–350). Amsterdam: North-Holland.
- Rochet, J.-C. (1987). A necessary and sufficient condition for rationalizability in the quasilinear context. *Journal of Mathematical Economics*, 16(2), 191–200.

- Rockafellar, R. T. (1966). Characterization of the subdifferentials of convex functions. *Pacific Journal of Mathematics*, 17(3), 497–510.
- Rockafellar, R. T. (1970). *Convex analysis*. Princeton, NJ: Princeton University Press.
- Saks, M. & Yu, L. (2005). Weak monotonicity suffices for truthfulness on convex domains. *Proceedings of the 6th ACM conference on electronic commerce*, 286–293.
- Sorin, S. (1999). Merging, reputation, and repeated games with incomplete information. *Games and Economic Behavior*, 29(1–2), 274–308.
- Sorin, S. (2002). *A first course on zero-sum repeated games*. Berlin: Springer.
- Spence, M. (1974). Competitive and optimal responses to signals: An analysis of efficiency and distribution. *Journal of Economic Theory*, 7(3), 296–332.
- Tarski, A. (1955). A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2), 285–309.
- Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operations Research*, 26(2), 305–321.
- Topkis, D. M. (1998). *Supermodularity and complementarity*. Princeton, NJ: Princeton University Press.
- van Zandt, T. & Vives, X. (2007). Monotone equilibria in Bayesian games of strategic complementarities. *Journal of Economic Theory*, 134(1), 339–360.
- Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19(3), 305–321.
- Vohra, R. V. (2011). *Mechanism design: A linear programming approach*. Cambridge: Cambridge University Press.
- Zhou, L. (1994). The set of Nash equilibria of a supermodular game is a complete lattice. *Games and Economic Behavior*, 7(2), 295–300.