Dominant Strategy Mechanism Design

Larry Blume

May 11, 2020

Introduction

The ideal kind of implementation is dominant strategy, for obvious reasons. A mechanism that implements an outcome in dominant strategies will be much more credible since dominant-strategy is a stronger solution concept than Nash equilibrium, and because it is "prior free"; its construction will not depend upon a particular set of information assumptions shared among the agents.

We will solve an implementation problem in dominant strategies to see how it can be done. The revelation principle still holds. Thus we can reduce the question to one of direct revelation games. We ask if truth-telling is a dominant strategy equilibrium.

The Assignment Problem

Recall the transferable utility assignment problem. Every individual announces his true preferences, and a linear program is solved for which maximizes total surplus. The set of dual solutions gives position prices. An optimal primal-dual pair (x, p) has a market equilibrium interpretation: If individuals are charged for their assignment according to p, then

- The individual assigned to position j and paying p_j will not be willing to pay p_k for any other position k;
- unmatched individuals receive 0 surplus;
- unmatched positions are priced at 0.

Implementation of the Optimal Assignment Correspondence

The optimal assignment problem maps each preference profile on $\mathcal{J}$ positions for $\mathcal{I}$ individuals into the set of optimal assignments. We would like to design a game for which truth-telling is a dominant strategy.

A special case is the problem of assigning a single object to one of a number of individuals. One game that solves the problem is the second price auction: Everyone announces a valuation for the object. The object is awarded to the bidder who announces the highest valuation, and she pays the second highest valuation. We have already seen that this is a market equilibrium for this allocation problem, and that announcing one's true valuation is a dominant strategy equilibrium for this game. Can this result be generalized to more than one object?

Assigning a Single Object

Suppose there is one object to be allocated to one of two individuals. $v_{11} = 2$ and $v_{21} = 1$. The optimal solution is obvious: Allocate the object to individual 1. The primal program is:

$$\begin{aligned} \max\ 2x_{11} + 1x_{21} & \\ \text{s.t.} \quad x_{11} &\leq 1, \\ x_{21} &\leq 1, \\ x_{11} + x_{21} &\leq 1, \\ x &\geq 0. \end{aligned}$$

The solution is obvious; $x_{11} = 1$ and $x_{21} = 0$.

The Dual

Review

The dual program is more interesting.

$$\begin{aligned} &\min\ s_1 + s_2 + q_1 \\ \text{s. t.}\quad & s_1 + q_1 \geq 2, \\ & s_2 + q_1 \geq 1, \\ & s_1, s_2, q_1 \geq 0. \end{aligned}$$

To find the dual solutions, observe that since the x_{21} constraint in the primal is slack, the corresponding shadow price s_2 must equal 0. Thus $q_1 \geq 1$. Solutions giving the value 2 are then $s_1 = 2 - q_1$, and the non-negativity constraint on s_1 implies that $q_1 \leq 2$. So the set of solutions to the dual are $\{(s_1, s_2, q_1) : s_1 = 2 - q_1, s_2 = 0, 1 \leq q_1 \leq 2\}$. This is the set of equilibrium surpluses and object rents.

The second-price auction outcome, $x_{11} = 1$ and $q_1 = 1$, is an optimal primal-dual pair. What does this suggest?

Direct Mechanisms

Suppose the following: A central administrator asks each individual i to submit a vector of valuations $\tilde{v}_i = (\tilde{v}_{i1}, \ldots, \tilde{v}_{iJ})$ for the j positions. The administrator computes an optimal allocation $\tilde{x}$ using the announced valuations $\tilde{v}_i$. He also computes a vector of position prices $\tilde{q} = (\tilde{q}_1, \ldots, \tilde{q}_J)$, one for each position, and charges the individual assigned to position j $\tilde{q}_j$. Individual i's utility from this procedure is then $\sum_j (v_{ij} - \tilde{q}_j)\tilde{x}_{ij}$.

This describes a direct revelation mechanism. Each player i has for actions vectors $v_i \in \mathbf{R}^{\mathrm{J}}$, announced types. A strategy maps true types into announced types. The payoff $u_i(\tilde{v}_i, \tilde{v}_{-i}, v_i)$ to individual i of announcing $\tilde{v}_i$ when others announce $\tilde{v}_{-i}$ and his true type is v_i, is $\sum_j (v_{ij} - \tilde{q}_j)\tilde{x}_j$, where $\tilde{x}$ and $\tilde{q}$ are a position assignment and payments to the center, computed by the center according to some allocation rule for mapping announced types into assignments and some payment function which maps announced types into transfers. Every such pair, allocation rule and transfer functions, determines a game.

Problem Formulation

Can an optimal assignment be realized as a DS equilibrium of a direct revelation game?

The market equilibrium condition, that no one wants to switch positions at the announced prices, is a necessary condition for incentive compatibility, so one might look to the primal and dual lps for a solution to this problem.

- In the single-object case the minimal dual price worked.
- In a multi-object game, the set of dual prices is a lattice, so there is a coordinatewise-minimal dual price.

This is a conjecture!

More Intuition

Denote by $V(g)$ the value function for the primal problem when there are g copies of the single object to allocate. The value function $V(g)$ is concave, and the dual prices comprise the set of supergradients of V at the point $g = 1$. The highest dual price, $\tilde{v}_1$, measures the welfare loss of decreasing the quantity available of the object, while $\tilde{v}_2$, the lowest prices, measures the welfare gain of increasing the quantity available of the object. If we could imagine removing the object, not allocating it, the lost welfare would be individual 1's, $\tilde{v}_1$. If a second copy of the object is made available, individual 1 is constrained to consume only 1, so the additional object would go to individual 2, who values it at rate $\tilde{v}_2$, and so the increase in announced welfare is $\tilde{v}_2$.

The Case of Multiple Positions

Suppose now that there are sets $\mathcal{I}$ of individuals and $\mathcal{J}$ of objects. We refer to both this problem and its value function by $V(\mathcal{I}, \mathcal{J})$.

Consider the algorithm which chooses the coordinate-wise minimum $\underline{q}$ and makes an optimal assignment based on the announced valuations. Call this algorithm the minimal shadow price mechanism.

Theorem. The minimal shadow-price mechanism is a truthful direct mechanism in dominant strategies for implementing the surplus-maximizing match.

Proof

Following the intuition of slide 9, we will consider prerturbations of $V(\mathcal{I}, \mathcal{J})$ that add and subtract individuals and positions.

The proof has three steps. Let $\iota : \mathcal{I} \to \mathcal{J}$ denote the optimal match.

1. If individual i is assigned to position j in the program $V(\mathcal{I}, \mathcal{J})$, then there is an optimal match to the program with one extra copy of position j in which she is also given position j.

2. Using 1, show that if $\iota(i) = j$, then

$$\underline{q}_j = V(\mathcal{I}/\{i\}, \mathcal{J}) - V(\mathcal{I}/\{i\}, \mathcal{J}/\{j\}).$$

   This is the key step: $\underline{q}_j$ does not depend on i's $\tilde{v}_i$.

3. Using 2, prove the theorem.

Proof of Step 3

Suppose that individuals 2 through N report vectors $\tilde{v}_i$, and that $\iota(1) = 1$. If individual 1 reports his true vector of valuations v_1, person 1 pays

$$V(\mathcal{I}/\{1\}, \mathcal{J}) - \sum_{i \neq 1} \tilde{v}_{i,\iota(i)}$$

which does not depend on what 1 announces. His gain is

$$v_{11} - \left(V(\mathcal{I}/\{1\}, \mathcal{J}) - \sum_{i \neq 1} \tilde{v}_{i,\iota(i)} \right) = V(\mathcal{I}, \mathcal{J}) - V(\mathcal{I}/\{1\}, \mathcal{J}).$$

This last term, bidder 1s marginal product, is the value of including individual 1 in the allocation process.

Proof of Step 3

If 1 announces a $\tilde{v}_1 \neq v_1$, there will be a new optimal allocation κ that maximizes $\sum_{i \neq 1} \sum_j \tilde{v}_{ij} x_{ij} + \sum_j \tilde{v}_{1j} x_{1j}$, and 1's payoff becomes

$$\sum_{i \neq 1} \tilde{v}_{i\kappa(i)} + v_{1\kappa(1)} - V(\mathcal{I}/\{1\}, \mathcal{J})$$

$$\leq \max \left\{ \sum_{i \neq 1} \sum_j \tilde{v}_{ij} x_{ij} + \sum_j v_{1j} x_{1j} \right\} - V(\mathcal{I}/\{1\}, \mathcal{J}).$$

$$= V(\mathcal{I}, \mathcal{J}) - V(\mathcal{I}/\{1\}, \mathcal{J}).$$

□

The idea that the payment does not depend upon 1's announcement is crucial here, as it is in the second-price auction.

Proof of Step 2

Optimal solutions to dual problems are supergradients for the primal value function. So for any $y \in \partial V(\mathcal{I}, \mathcal{J})$ and position j,

$$V(\mathcal{I}, \mathcal{J} \cup \{j\}) \leq V(\mathcal{I}, \mathcal{J}) + y_j.$$

Because the value function is piecewise linear, this holds for equality for some supergradient q, and concavity also implies that this must be the smallest magnitude directional derivative in the direction j, which is $\underline{q}_j$.

Proof of Step 2

Optimal solutions to dual problems are supergradients for the primal value function. So for any $y \in \partial V(\mathcal{I}, \mathcal{J})$ and position j,

$$V(\mathcal{I}, \mathcal{J} \cup \{j\}) \leq V(\mathcal{I}, \mathcal{J}) + y_j.$$

Because the value function is piecewise linear, this holds for equality for some supergradient q, and concavity also implies that this must be the smallest magnitude directional derivative in the direction j, which is $\underline{q}_j$.

This seems to use vector space reasoning, but we don't have a vector space. But we do! Represent any set S by the vector e^S such that $e^S_k = 1$ if $k \in S$ and 0 otherwise. Now V is defined on a vector space. A participation constraint for person i is that $\sum_j x_{ij} \leq \sum_k e^{\{i\}}_k$, etc. Adding a bunch of f_1 persons of type 1, f_2 persons of type 2, etc.m just changes the participation constraint for i to $\sum_j x_{ij} \leq \sum_k e^{\{i\}}_k + f_i$. This is a linear change in constraints, and the usual rules of differentiation apply.

Proof of Step 2

Assume that person 1 is optimally assigned to position 1. Then from step 1, $V(\mathcal{I}, \mathcal{J} \cup \{1\}) = v_{11} + V(\mathcal{I}/\{1\}, \mathcal{J})$. Therefore

$$
\begin{aligned}
\underline{q}_1 &= V(\mathcal{I}, \mathcal{J} \cup \{1\}) - V(\mathcal{I}, \mathcal{J}) \\
&= (v_{11} + V(\mathcal{I}/\{1\}, \mathcal{J})) - (v_{11} + V(\mathcal{I}/\{1\}, \mathcal{J}/\{1\})) \\
&= V(\mathcal{I}/\{1\}, \mathcal{J}) - V(\mathcal{I}/\{1\}, \mathcal{J}/\{1\})
\end{aligned}
$$

□

Proof of Step 1

Call the problem with sets $\mathcal{I}$ and $\mathcal{J}$ the "old problem" and the problem with sets $\mathcal{I}$ and $\mathcal{J} \cup \{1\}$ the "new problem". The hypothesis is that in the old problem 1 is assigned to 1, and the claim is that the new problem has an optimal solution wherein 1 is assigned to 1.

In the new problem, if at least one of the two 1 positions is empty, there is no cost to removing it, the old and new problems are then the same, and so the old problem optimal solution also solves the new.

So suppose the new problem has an optimal solution in which both 1 positions are assigned, but neither to person 1.

Proof of Step 1

We divide the individuals into three groups depending on their assignment in the new problem. $\mathcal{I}_0$ is the set of people whose assignment is the same in the old problem and in the new. Now choose one of the individuals assigned to a position 1, say i_1. He left his old position, and it was assigned to i_2. She left her old position, and it was assigned to i_3, etc. This chain ends in one of two ways: Either it comes around to person 1, who gave up his old position and it went to i_1, or it reaches a position that in the new problem's optimal solution is not assigned. If we construct two chains starting from the two individuals assigned to the two position 1s, one will end the first way; call the group of people in that chain $\mathcal{I}_2$. And call the group of people in the remaining chain $\mathcal{I}_3$.

Proof of Step 1

The value of the new solution is the sum of the values of each groups assignments. The value of group $\mathcal{I}_0$ assignments is the same in both problems because these assignments are unchanged. Group I_1 just swapped their old positions with each other.

The value of this group's new assignments must equal the value of their old assignments. If the new assignments for this group have have higher total value than do the old assignments, then the old assignment could not be optimal in the old problem. Why? Group $\mathcal{I}_1$ has only swapped amongst themselves. Leaving the old group $\mathcal{I}_0$ and $\mathcal{I}_2$ assignments unchanged in the old problem, and giving $\mathcal{I}_1$ the new assignments creates a feasible assignment for the old problem with higher value, contradicting the optimality of the old solution for the old problem. A similar argument shows that the new assignments cannot be optimal in the new problem if the $\mathcal{I}_1$ value is lower.

Proof of Step 1

To conclude the proof simply note that if one replaces the new $\mathcal{I}_1$ assignments with the old $\mathcal{I}_1$ assignments, this allocation is also optimal. And, person 1 is in group I_1. □

Dominant Strategy Equilibrium

Review

A strategy σ_i^* for individual i is always optimal iff for all strategy profiles σ_{-i}, strategies σ_i, and type profiles v,

$$u_i\big(\sigma_i^*(v_i), \sigma_{-i}(v_{-i}, v_i)\big) \geq u_i\big(\sigma_i(v_i), \sigma_{-i}(v_{-i}, v_i)\big).$$

It is dominant if, in addition, for every $\sigma_i \neq \sigma_i^*$ there is a type profile v such that the preceding inequality is strict.

This is a "worst-case" analysis of sorts. A dominant strategy requires that an individual do well no matter what the types. A more relaxed Bayesian alternative would require doing well only on average with respect to the common prior. Of course multiple prior models allow us to imagine all kinds of in-between cases.

TU Dominant Strategy Implementation

A social allocation problem is described by a group of I individuals, and a set C of social choices. Each individual i is described by a set V_i of types. Let $V = \prod_i V_i$. Each individual i's preferences over social decisions is determined by her type, and represented by a utility function $u_i(c, v_i)$. Individuals have money holdings too, that enter into utilities. In order to provide incentives for efficient behavior, individuals may be required to pay a tax or receive a subsidy. We assume that all individuals are sufficiently wealthy to afford any transfers the central administrator might impose. (Otherwise utility might not be transferable.) Thus i's utility from the social choice c and holding money e_i when her type is v_i, is $U_i(c, e_i, v_i) = u_i(c, v_i) + e_i$.

TU Dominant Strategy Implementation

The prescribed choices to implement are described by an allocation rule:

An allocation rule is a function $r : V \to C$ that assigns a social choice c to each vector $v \in V$ of types.

Much of the mechanism design literature is concerned with implementing efficient allocation rules.

An allocation rule r is efficient iff $\sum_i u_i\big(r(v), v_i\big) \geq \sum_i u_i(c, v_i)$ for all $v \in V$ and $c \in C$.

Given transferable utility, this definition of efficiency is equivalent to Pareto optimality of the social decision.

TU Dominant Strategy Implementation

The prescribed choices to implement are described by an allocation rule:

An allocation rule is a function $r : V \to C$ that assigns a social choice c to each vector $v \in V$ of types.

Much of the mechanism design literature is concerned with implementing efficient allocation rules.

An allocation rule r is efficient iff $\sum_i u_i(r(v), v_i) \geq \sum_i u_i(c, v_i)$ for all $v \in V$ and $c \in C$.

Given transferable utility, this definition of efficiency is equivalent to Pareto optimality of the social decision.

Although most of mechanism design is about the implementation of efficient decision rules, it need not be. (There is a literature on the characterization of those properties of potential decision rules that make them implementable. See Maskin TK.)

TU Dominant Strategy Implementation

In a direct mechanism individuals announce a type, $\tilde{v}_i$. The planner maps the announced type into an outcome; thus the planner's action can be described by an allocation rule $\tilde{r}$. In order to control individuals' behaviors, taxes and subsidies may be necessary. So in addition to announcing a social choice, the planner also announces a tax or subsidy from or to each individual.

A transfer is a map $t : V \to \mathbf{R}^I$, which describes how much an individual needs to pay if the announced types are v. A transfer is feasible if $\sum_i t_i(v) \geq 0$ for all $v \in V$. A direct mechanism is a pair $\langle \tilde{r}, \tilde{t} \rangle$ where $\tilde{r}$ is an allocation rule and $\tilde{t}$ is a feasible transfer.

TU Dominant Strategy Implementation

In a direct mechanism individuals announce a type, $\bar{v}_i$. The planner maps the announced type into an outcome; thus the planner's action can be described by an allocation rule $\bar{r}$. In order to control individuals' behaviors, taxes and subsidies may be necessary. So in addition to announcing a social choice, the planner also announces a tax or subsidy from or to each individual.

A transfer is a map $t : V \to \mathbf{R}^I$, which describes how much an individual needs to pay if the announced types are v. A transfer is feasible if $\sum_i t_i(v) \geq 0$ for all $v \in V$. A direct mechanism is a pair $(\bar{r}, \bar{t})$ where $\bar{r}$ is an allocation rule and $\bar{t}$ is a feasible transfer.

There is also a literature on implementary without monetary transfers — for example, school matching.

TU Dominant Strategy Implementation

A social choice rule r is implemented in always optimal (dominant) strategies by the mechanism $\langle M, \hat{r}, \hat{t} \rangle$ if for each i there exists a function $\sigma_i : V_i \to M_i$ such that

- $\sigma_i(v_i)$ is an always optimal (dominant) strategy;
- for all $v \in V$, $\hat{r}(\sigma(v)) = r(v)$.

A social choice rule that can be implemented by always optimal strategies in some mechanism is incentive-compatible or strategy-proof.

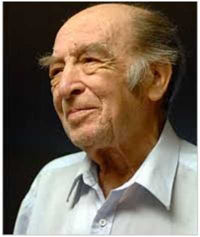

Mechanism is the key concept in implementation theory and market design. The concept was introduced by Leo Hurwicz (1960) who later received a Nobel prize for his work.

TU Dominant Strategy Implementation

A social choice rule r is implemented in always optimal (dominant) strategies by the mechanism $\langle M, \tilde{x}, \tilde{t} \rangle$ if for each i there exists a function $\sigma_i : V_i \to M_i$ such that

- $\sigma_i(v_i)$ is an always optimal (dominant) strategy;
- for all $v \in V$, $\tilde{x}(\sigma(v)) = r(v)$.

A social choice rule that can be implemented by always optimal strategies in some mechanism is incentive-compatible or strategy-proof.

Mechanism is the key concept in implementation theory and market design. The concept was introduced by Leo Hurwicz (1960) who later received a Nobel prize for his work.

"Optimality and informational efficiency in resource allocation processes", in Arrow, Karlin and Suppes (eds.), *Mathematical Methods in the Social Sciences.* Stanford University Press. 1960. Also read: Hurwicz, L. "The design of mechanisms for resource allocation", *American Economic Review* 63 (1973), Papers and Proceedings, 1–30.

The difference between always optimal and dominant is that if types are sufficiently restricted, a particular type may never have an impact on the outcome. In that case, the actions of such a type could never be dominant. When there are enough types, this won't happen.

TU Dominant Strategy Implementation

If a mechanism $\langle M, \hat{r}, \hat{t} \rangle$ implements a social choice rule r with dominant strategy profile σ, then there is a direct mechanism with allocation rule r, $\langle V, r, t \rangle$ such that

- Truth-telling is a dominant strategy; and
- $t(v) = \hat{t}\big(\sigma(v)\big)$ for all $v \in V$.

Proof. If this were false, then for some individual i, type profile v' and announced type $\tilde{v}_i$,

$$\tilde{U}_i(\tilde{v}_i, v'_{-i}, v_i) > \tilde{U}_i(v'_i, v'_{-i}, v'_i)$$

that is

$$\hat{U}\big(\sigma_i(\tilde{v}_i), \sigma_{-i}(v'_{-i}), v'_i\big) > \hat{U}\big(\sigma_i(v'_i), \sigma_{-i}(v'_{-i}), v'_i\big).$$

Thus the strategy which plays $\hat{\sigma}_i(v_i) = \sigma_i(v_i)$ when $v_i \neq v'_i$ and $\hat{\sigma}_i(v_i) = \sigma_i(\tilde{v}_i)$ when $v_i = v'_i$ performs better than does σ_i when the type vector of is v'; thus σ_i is not a dominant strategy. □

VCG Mechanisms

A direct mechanism $\langle \tilde{r}, \tilde{t} \rangle$ is a VCG mechanism iff

- $\tilde{r}$ is an efficient decision rule; that is, it maximizes $\sum_i v_i(c, \tilde{v}_i)$, and
- if for every $i \in I$ there is a function $\tau_i : V_{-i} \to \mathbf{R}$ such that for all $v \in V$,

$$\tilde{t}_i(v) = \tau_i(v_{-i}) - \sum_{j \neq i} v_j\big(\tilde{r}(v), v_j\big).$$

Pivot Mechanism

The *pivot mechanism* is the VCG mechanism in which each $\tau_i(v_{-i}) \equiv \max_c \sum_{j \neq i} u_j(c, v_{-i})$. The idea behind the pivot mechanism is that individual i pays his marginal contribution to the social welfare of others: What they would have achieved without him less what they achieve with him present. The minimum-support-price mechanism is a pivot mechanism for the assignment problem. Notice too that an individual who has no effect on the allocation pays 0.

VCG and Implementation

Theorem. VCG mechanisms are dominant-strategy incentive compatible.

Proof. For any VCG mechanism,

$$\tilde{U}_i(\tilde{v}_i, \tilde{v}_{-i}, v_i) = \sum_j u_j\big(\tilde{r}(\tilde{v}), v_i\big) - \tau_i(\tilde{v}_{-i}).$$

The optimal strategy for individual i is to choose v_i to maximize social welfare given any v_{-i}. □

The converse is true to under some additional hypotheses, but it is not easy to show.