

Exercises for Lecture 2 (M1): Revelation Principle, DSIC Mechanisms

Problem 1: Søndre campus

There are currently talks at KU about moving the Faculty of Social Sciences from the Kommunehospitalet that we occupy now to Søndre campus, where some other faculties are currently located.¹ The costs and benefits of such a move are currently being evaluated. Some, however, see this whole discussion as a bargaining maneuver in the upcoming negotiations with the firm that owns the Kommunehospitalet and leases it to the university – a credible threat of leaving may help the university bargain a better lease rate.

Your mission is to frame this choice of whether SAMF should move as a mechanism design problem. The goal of the mechanism is to extract the information about costs and benefits of the potential move from the relevant parties. In particular, answer the following questions within this setting:

1. Who is the designer?
2. What is the outcome in this setting? (Do we have access to transfers? Is the set of allocations k given by simply $K = \{\text{move}, \text{no move}\}$ or is it more multifaceted?)
3. Who are the players?
4. What information do the players have that is relevant to determining the optimal outcome/allocation?
5. How would you model the players' utility functions? (Give a concrete example.)
6. What criteria or conditions should the mechanism satisfy?
7. What would be the desirable outcome/allocation rule that you would want to implement with such a mechanism? How can you check whether this rule is, in fact, implementable?
8. If you allowed for transfers: how would you proceed with designing transfers that support the chosen allocation rule? (You do not need to actually derive the transfers.)
9. How would your mechanism work in the real world, in terms of organization and logistics?²

NOTE: treat this as a real-life assignment from the university officials. Your goal is to give the best possible answer to the question they ask, NOT to frame the problem in the simplest way possible. That said, you should still be realistic and try to set up the problem in a way that would be tractable and doable given the resources available to a committee responsible for this decision.

Solution

This is an open-ended question, so many answers are possible. Below is one example.

1. The designer is the university leadership (rector), possibly proxied by a “committee on moving”.
2. It feels somewhat strange to include monetary transfers in this problem. While it may be fine to pay small amounts to students and faculty for completing a survey or pay departments to compile a report on a given issue, making these payments contingent on responses to provide incentives is the weird part, which will likely not be well accepted. Therefore, an outcome is simply an allocation.

¹News article from Uniavisen (in Danish): <https://tinyurl.com/y4uwrefe>.

²Example: “all faculty, staff, and students must post a note on the door of their office which would contain their report of something; a dedicated person will walk around and enter responses in an excel sheet, which will then be used to determine the outcome”.

An allocation, however, is much richer than just a binary decision. In case of a move we would also need to decide how to allocate the spaces on Søndre campus between the departments and faculties, whether to build new lecture halls (or force the students and faculty to commute to CSS or Nørre campus for classes), whether and how to merge the duplicating departments, etc.

3. The set of players should include anyone who possesses information relevant to the outcome, and has preferences regarding the outcome, which could prevent them from communicating this information to us truthfully. In this problem, this includes:

- Students, faculty, and staff, who all have private valuations for the move.
- Future students, for the same reason. We obviously cannot include future students in our mechanism because we do not know who they are, so the next best option is to let the current students speak on their behalf.
- The university's building administration, which knows how much room capacity SAMF requires, how much capacity is available on Søndre campus, how much the exploitation of these rooms costs on both campuses, and how much it would cost to build more lecture halls on Søndre campus. It is not immediate that there exists any conflict preventing this information from being openly communicated, but there could be some. E.g., I believe that at the moment, these administrations exist separately on different campuses, so they may be opposed to merging or, conversely, they may actively prefer the merge.
- Some university departments (e.g., IT) can better estimate the cost of the actual process of moving. Potential conflicts of interest are as above.
- The firm which owns the Kommunehospitalet and leases it to the university – it likely has some understanding of what its outside option is in case the university leaves, and the company would have to find new tenants for this property. Note that it makes a lot of sense to allow monetary transfers with this particular player.
- ...

4. See above.

5. For concreteness, let us suppose from this point onwards that the costs of the move can be evaluated by internal departments without any conflict and need for a mechanism. Then our set of players is narrowed down to two groups: stakeholders (students, faculty, and staff) and the propertyowner firm. The firm's utility can be modelled as:

$$u_f(x, \theta) = \begin{cases} -t_f & \text{if no move;} \\ \theta_f & \text{if move,} \end{cases}$$

where again the firm's type θ_f is its outside option relative to continuing the current agreement, and t_f is the negative of the change in lease that KU pays for the Kommunehospitalet. The firm does not care about the details of the move if it happens.

With stakeholders it is a little more difficult, since the details of the allocation now matter. What we can do is assume that the allocation can be split into a number of aspects $l \in \{1, \dots, L\}$ and each k_l can be represented as a number. E.g., one aspect is "will all classes be held in the same place after the move", another is "if yes, will this place be Søndre campus", another is "will there be place for a student bar", etc. If we take this approach, then stakeholder i 's type can be represented as a vector of valuations for every aspect $\theta_i = (\theta_{i,1}, \theta_{i,2}, \dots, \theta_{i,L})$, and then the utility can be approximated as $u_i(x, \theta) = \sum_{l=1}^L k_l \theta_{i,l}$.

6. We obviously want the mechanism to be incentive compatible, at least in the Bayesian sense.³ Indi-

³Large number of players is one excuse to switch from DSIC to BIC: heuristically, the more possible type profiles θ_{-i} other players collectively have, the less important every single individual case is for i 's expected utility at the point where they are

vidual rationality is not an issue, once you realize that “not responding to a survey” is just a kind of response. While students and staff and faculty have an outside option of leaving KU and applying to another university, this is likely a costly option. You can, however, make a case for IR being a desirable condition to satisfy for future students (so they choose KU over other universities), as well as the firm that owns Kommunehospitalet (depending on how you model its utility function). Finally, since transfers are not a part of the problem (except when dealing with the firm), budget balance is not a particularly relevant requirement.

7. The simplest answer: the desired allocation rule $k^d(\theta)$ should maximize the stakeholders' welfare, i.e., the sum of utilities net of the costs borne by the university. Note that this is not the efficient allocation according to the standard definition, since it ignores the firm's well-being. You can also present various social choice arguments saying that the sum of utilities is not the best measure to use here. E.g., depending on the university's priorities, you may want to assign different weights to students' well-being versus that of the faculty.

After calculating this allocation rule, you can realize that this problem fits the general setting (but not the quasilinear or Euclidean, since we do not allow for transfers), so you can use weak monotonicity of the outcome function/allocation rule to test for implementability. (Again, weak monotonicity for the general setting, not the weak monotonicity for the Euclidean setting, even though k is a vector of numbers!)

8. I went with a mix, allowing transfers to/from the firm, but not the stakeholders. If the firm's utility entered the objective function, we could use the (g)VCG transfers to align the firm's objective with the desired objective. However, the firm's utility is not a part of the objective function, so the VCG does not work as is. There is no set recipe for this case, but you can employ the first principles to come up with a mechanism that may not necessarily be best, but is good enough. In particular, the firm must be unable to change its transfer without also altering the allocation rule.

The simplest (but not the only!) way to achieve this is to give the firm no power over rent. E.g., the university can determine its willingness to pay for the Kommunehospitalet given stakeholders' reports, and then present the firm with a take-it-or-leave-it offer.

9. It is probably easiest to set up an online survey for the stakeholders, and no issues should arise there. Negotiations with the firm are a more subtle matter, since using a direct revelation mechanism may not be appealing for the firm if it does not believe in the university's power to commit to the mechanism, or if it thinks its report may be used against it in future negotiations. A take-it-or-leave-it-offer as suggested above solves this issue.

Problem 2: Screening

One application of mechanism design is to profit maximization when consumers have private information about their valuations. One example of such a problem is second-degree price discrimination that you have seen in mikro II. The following is a variation of that, known as a “monopolistic screening” problem (with two types) that you may have seen in game theory.

Suppose a seller-designer offers a single product for sale that he can produce at zero cost. She offers a menu of pairs of quantities $k \in [0, 1]$ and payments $t \in \mathbb{R}_+$ (for the whole amount k , not per unit). There is one buyer with valuation $\theta \in \{L, H\}$ for the product, which is his private information. The seller's belief regarding θ is given by $\phi(H) = \phi$, $\phi(L) = 1 - \phi$. The buyer's preferences are given by $u_b(k, t, \theta) = \theta k - t$ if he buys the product and zero otherwise.

1. Explain why it is sufficient for the seller to offer a menu consisting of two items: (k_H, t_H) and (k_L, t_L) .

making a decision.

2. Write down the seller's problem of maximizing her expected profit subject to the buyer's incentive compatibility (IC) and individual rationality (IR) constraints for every θ , in terms of the model primitives.
3. Derive the seller's optimal menu $((k_H^*, t_H^*), (k_L^*, t_L^*))$ by following the steps below.
 - (a) Show that if $((k_H, t_H), (k_L, t_L))$ satisfy IC_H and IR_L , then they also satisfy IR_H .
 - (b) Show that $((k_H, t_H), (k_L, t_L))$ satisfy IC_H and IC_L only if $k_H \geq k_L$.
 - (c) Show that if $((k_H, t_H), (k_L, t_L))$ are such that $k_H \geq k_L$ and IC_H binds (i.e., is satisfied with equality), then they also satisfy IC_L .
 - (d) Show that given all of the above, it is always optimal to choose $((k_H, t_H), (k_L, t_L))$ in such a way that IC_H and IR_L bind.
 - (e) Given all of the above, solve for the optimal menu $((k_H^*, t_H^*), (k_L^*, t_L^*))$.

Solution

1. By the revelation principle, the problem can be solved by a direct revelation mechanism that asks the consumer for his type and prescribes an outcome (k, t) in return. Since there are only two types of the consumer, such a DRM will respond with either (k_H, t_H) or (k_L, t_L) , depending on the consumer's report of θ . An equivalent (indirect) mechanism would, instead of asking the consumer to report his type, simply offer him a choice between these two outcomes.
2. The problem looks as follows:

$$\begin{aligned} & \max_{((k_H, t_H), (k_L, t_L))} \{\phi t_H + (1 - \phi)t_L\} \\ \text{s.t. } & (IC_H): \theta_H k_H - t_H \geq \theta_H k_L - t_L, \\ & (IC_L): \theta_L k_H - t_H \leq \theta_L k_L - t_L, \\ & (IR_H): \theta_H k_H - t_H \geq 0, \\ & (IR_L): \theta_L k_L - t_L \geq 0. \end{aligned}$$

3. Proceeding step by step:
 - (a) Consider the following chain of inequalities:

$$\theta_H k_H - t_H \geq \theta_H k_L - t_L \geq \theta_L k_L - t_L \geq 0.$$

The first inequality above is IC_H , the second is satisfied since $\theta_H > \theta_L$, and the last one is IR_L . Hence if IC_H and IR_L hold, then $\theta_H k_H \geq 0$, meaning that IR_H holds automatically, and can be ignored.

- (b) Subtracting IC_L from IC_H , we get

$$\begin{aligned} \theta_H k_H - t_H - (\theta_L k_H - t_H) & \geq \theta_H k_L - t_L - (\theta_L k_L - t_L) \\ \iff (\theta_H - \theta_L)k_H & \geq (\theta_H - \theta_L)k_L \\ \iff (\theta_H - \theta_L)(k_H - k_L) & \geq 0. \end{aligned}$$

Since $\theta_H > \theta_L$, we get that IC_H and IC_L together imply that $k_H \geq k_L$.

(c) If IC_H binds, then $t_H - t_L = (k_H - k_L)\theta_H$. Plugging that into IC_L gives:

$$\begin{aligned} \theta_L k_H - t_H &\leq \theta_L k_L - t_L \\ \iff \theta_L(k_H - k_L) &\leq t_H - t_L \\ \iff \theta_L(k_H - k_L) &\leq (k_H - k_L)\theta_H \\ \iff (\theta_H - \theta_L)(k_H - k_L) &\geq 0. \end{aligned}$$

(Same expression as above, but a slightly different implication behind it. Previously the idea was that this condition was necessary for IC_H and IC_L to hold. Now the idea is if IC_H holds with equality, then the condition above is equivalent to IC_L .)

Since $\theta_H > \theta_L$ and it is assumed that $k_H \geq k_L$, the condition above holds, hence IC_L also holds.

(d) In part (a) we saw that IR_H can be ignored because it is always satisfied by a menu that satisfies the other constraints. In part (c) we saw that if the optimal menu is such that IC_H binds, then IC_L can also be ignored. Let us consider a problem without IR_H and IC_L , and we will show that IC_H does indeed bind in the optimum, hence IC_L is indeed satisfied.

$$\begin{aligned} \max_{((k_H, t_H), (k_L, t_L))} & \{ \phi t_H + (1 - \phi) t_L \} \\ \text{s.t. } (IC_H) : & \theta_H k_H - t_H \geq \theta_H k_L - t_L, \\ (IR_L) : & \theta_L k_L - t_L \geq 0. \end{aligned}$$

Proceed by contradiction: suppose IC_H is slack in the optimum (i.e., the optimal menu is such that IC_H holds with strict inequality). Then the seller could offer a different menu with a slightly higher t_H – this would yield higher expected profit, and both constraints would still be satisfied, hence such a menu is strictly better, which contradicts the original menu being optimal. Therefore, IC_H can not be slack in the optimal menu. Similarly, if IR_L is slack in the optimum, then increasing t_L and t_H is a similar profitable deviation, yielding another contradiction.

(e) Given all of the above, the problem can be re-stated as

$$\begin{aligned} \max_{((k_H, t_H), (k_L, t_L))} & \{ \phi t_H + (1 - \phi) t_L \} \\ \text{s.t. } (IC_H) : & \theta_H k_H - t_H = \theta_H k_L - t_L, \\ (IR_L) : & \theta_L k_L - t_L = 0, \\ & k_H \geq k_L. \end{aligned}$$

This problem is much simpler than the original one. From IR_L , we get $t_L = \theta_L k_L$, and then from IC_H , we can express $t_H = \theta_H(k_H - k_L) + \theta_L k_L$. The problem is then equivalent to

$$\begin{aligned} \max_{(k_H, k_L)} & \{ \phi \theta_H (k_H - k_L) + \theta_L k_L \}, \\ \text{s.t. } & k_H \geq k_L. \end{aligned}$$

This is a linear maximization problem with linear constraints (remember also that $k_\theta \in [0, 1]$), hence we'll have a corner solution. There are three corners in the constrained set: $(k_H, k_L) \in \{(0, 0), (1, 0), (1, 1)\}$. It is easy to see that $(0, 0)$ is dominated by $(1, 0)$. Either of the other two

can be optimal depending on the parameters. Hence the optimal menu is given by:

$$((k_H^*, t_H^*), (k_L^*, t_L^*)) = \begin{cases} ((1, \theta_H), (0, 0)) & \text{if } \phi\theta_H \geq \theta_L, \\ ((1, \theta_L), (1, \theta_L)) & \text{if } \phi\theta_H \leq \theta_L. \end{cases}$$

Problem 3: Screening 2

This is a marginally more difficult version of the previous problem. Once you understood the solution of the previous problem, try to solve this one by following the same algorithm.

The Chicago Transit Authority (the organization in charge of the Chicago subway system) has decided that it needs to do more to maximize its revenue. As such it has hired you to design its new price and service scheme. There are two types of customers, High-class and Low-class. They have preferences over the fare P and the degree of bad smell in the train car they ride in, denoted by B . They have told you that they are able to charge different fares depending on the car a customer rides in (i.e., to have different classes of service).

The type of a customer is not observable; the fraction of high-class customers is λ . Customers' utility functions are $u_i(P, B) = v - \theta_i P - B$, for $i = H, L$, where $\theta_L > \theta_H > 0$. All customers get utility (normalized) of 0 from walking (their next best alternative) instead of taking the CTA train.

Making train cars smell bad is not costless (workers need to be hired to rub garbage on the seats): the CTA incurs a cost of $\gamma B > 0$ per customer who rides in a car that has smell level B .

1. Write down the problem you would solve for determining the CTA's profit-maximizing scheme. Assume throughout that the CTA cannot charge negative prices; i.e., that $P \geq 0$. Assume also that the CTA wants to serve both high and low class customers.
2. Determine the CTA's profit-maximizing scheme. How does it depend on the parameters of the problem?

Solution

1. Appealing to the Revelation Principle, the CTA's problem can be written as:

$$\begin{aligned} \max_{(P_L, B_L) \geq 0, (P_H, B_H) \geq 0} & \left\{ \lambda(P_H - \gamma B_H) + (1 - \lambda)(P_L - \gamma B_L) \right\} \\ \text{s.t. } (IR_L) & \quad v - \theta_L P_L - B_L \geq 0 \\ (IR_H) & \quad v - \theta_H P_H - B_H \geq 0 \\ (IC_L) & \quad v - \theta_L P_L - B_L \geq v - \theta_L P_H - B_H \\ (IC_H) & \quad v - \theta_H P_H - B_H \geq v - \theta_H P_L - B_L. \end{aligned}$$

2. Observe first that (IR_H) is redundant since it is implied by (IC_H) plus (IR_L) .

Observe next that (IR_L) will bind: otherwise we could raise both P_L and P_H by some $\varepsilon > 0$ and not violate any constraints, thereby raising profit.

Observe next that $B_H = 0$. If not, then we can lower B_H and raise P_H so that $\theta_H P_H + B_H$ is unchanged. This violates no constraints (since it raises $\theta_L P_H + B_H$), but raises profit.

Let's ignore (IC_L) and check that it holds at the end. In this case, we must have (IC_H) binding: otherwise we could lower B_L a little and raise profit. Hence, we see that (by rewriting the binding

IR_L and IC_H respectively)

$$P_L = \frac{v - B_L}{\theta_L}$$

$$P_H = P_L + \frac{B_L}{\theta_H} = \frac{v - B_L}{\theta_L} + \frac{B_L}{\theta_H}$$

So the CTA should choose $B_L \in [0, v]$ to solve

$$\max_{B_L} \lambda \left(\frac{v - B_L}{\theta_L} + \frac{B_L}{\theta_H} \right) + (1 - \lambda) \left(\frac{v - B_L}{\theta_L} - \gamma B_L \right).$$

This problem is linear, so the solution is to set $B_L = v$ if

$$\left(\frac{\lambda}{1 - \lambda} \right) \left(\frac{1}{\theta_H} - \frac{1}{\theta_L} \right) - \left(\frac{1}{\theta_L} + \gamma \right) > 0,$$

and set $B_L = 0$ if this inequality is reversed. Note that in the latter case, we have $P_L = P_H = \frac{v}{\theta_L}$, while in the former case we have $P_H - P_L > 0$. (And don't forget to verify that IC_L holds in both cases.) Also, observe that we are more likely to set $B_L = v$ when the fraction of high class consumers (λ) is high and the cost of making train cars smelly (γ) is low.

Problem 4: Judicial design

A suspect is in custody, accused of murder. If he goes to trial he will either be convicted or acquitted. If he is convicted he will be sent to prison for life giving him a payoff of -1 . If he is acquitted he goes free and has a payoff of 0 . The district attorney can offer plea bargains: allowing the defendant to plead guilty in return for a lighter sentence. In particular, for any $r \in (0, 1)$, the DA can offer a reduced sentence which, if accepted, would give the defendant a payoff of $-r$.

The defendant is privately informed about his chances for acquittal at trial: $\theta \in [0, 1]$ is the defendant's privately known probability of acquittal. If the defendant does not enter into a plea bargain with the DA he will go to trial and be convicted with probability $1 - \theta$.

Consider the mechanism design problem where the DA is the principal and the defendant is the agent. A social choice function is a mapping $f : [0, 1] \rightarrow \{\text{trial}\} \cup (0, 1)$ where $f(\theta) = \text{trial}$ means that type θ will go to trial and $f(\theta) = r \in (0, 1)$ means that type θ accepts a plea bargain giving him a sentence with payoff $-r$. DA thinks θ has full support on $[0, 1]$.

1. Write down the inequalities that characterize whether some given social choice function f is incentive-compatible for the defendant.
2. What is the set of all incentive-compatible social choice functions? You can proceed in the following steps:
 - Show that in any IC f at most one plea bargain r is available.
 - Show that f must be of cutoff type, with the suspect taking the plea if $\theta < \bar{\theta}$ and going to court otherwise.
 - Find the value of r that makes the cutoff s.c.f. f incentive compatible given some cutoff type $\bar{\theta}$.
 - Combine all of the above to characterize the set of implementable f .

Solution

By going to trial a defendant of type θ receives (expected) utility of $-(1 - \theta)$, while from accepting a plea bargain his utility is $-r$. Fix some s.c.f. $f(\theta)$. Let Θ_p be the set of types who are offered a plea bargain $f(\theta) = r(\theta)$, and Θ_t be the set of types who are meant to go to trial: $f(\theta) = \text{trial}$ ($\Theta_t \cup \Theta_p = [0, 1]$). Then the IC constraints are given by:

$$\begin{aligned} \text{for all } \theta \in \Theta_p : & \quad -r(\theta) \geq -r(\theta') \text{ for all } \theta' \in \Theta_p \\ & \quad \text{and } -r(\theta) \geq -(1 - \theta); \\ \text{for all } \theta \in \Theta_t : & \quad -(1 - \theta) \geq -r(\theta') \text{ for all } \theta' \in \Theta_p. \end{aligned}$$

We will characterize the set of IC social choice functions by a series of claims.

claim 1 $f(\cdot)$ has at most one value on the real line.

Proof: if $f(\theta_1) < f(\theta_2)$, $\theta_1, \theta_2 \in [0, 1]$ then a defendant of type θ_2 gains higher utility by declaring θ_1 (as $-f(\theta_1) > -f(\theta_2)$). This implies the mechanism is not IC for θ_2 .

claim 2 $f(\cdot)$ has a cutoff at some $\bar{\theta}$. i.e. $f(\theta) = \begin{cases} r & \text{if } \theta < \bar{\theta} \\ T & \text{if } \theta \geq \bar{\theta} \end{cases}$ (value at $\bar{\theta}$ is not unique)

Proof: assume $\theta' > \theta$, $f(\theta) = T$, $f(\theta') = r$. By IC for θ we know that $-r \leq -(1 - \theta)$. However, as $-(1 - \theta') > -(1 - \theta)$, this implies that $-(1 - \theta') > -r$ and we don't have IC for θ' .

claim 3 $r = 1 - \bar{\theta}$

Proof: $r \leq 1 - \bar{\theta}$ follows directly from IC for type $\bar{\theta}$, while $r \geq 1 - \bar{\theta}$ follows from IC of type $\bar{\theta} + \epsilon$. If type $\bar{\theta}$ were strictly better off by accepting the plea bargain, by continuity and monotonicity of benefit of trial, type $\bar{\theta} + \epsilon$ would also strictly prefer the plea bargain, contradicting IC for that type.

These four claims imply that for any $(r, \bar{\theta})$ s.t. $r = 1 - \bar{\theta}$ the social choice function

$$f(\theta) = \begin{cases} r & \text{if } \theta < \bar{\theta} \\ T & \text{if } \theta \geq \bar{\theta} \end{cases}$$

is incentive compatible.

Problem 5: Second-price auction

There is a single item being sold via a second-price sealed bid auction. There are $i = 1, \dots, n$ bidders. Every bidder i has a private valuation θ_i , which the other players believe is distributed according to some c.d.f. $F_i(\theta_i)$. All bidders simultaneously submit bids b_i to the seller (without seeing what the others bid). The highest bidder then wins the object and pays the second-highest bid $b_{(2)}$, their utility is then given by $u_i = \theta_i - b_{(2)}$. All other bidders get nothing and pay nothing, so their utility is zero.

1. Show that bidding truthfully ($b_i = \theta_i$) is a weakly dominant strategy for every bidder i .
2. Conclude that the second-price auction implements the efficient allocation rule in dominant strategies.

5.1 Solution

Part 1. Let's prove that for any θ_i and for any b_{-i} bidding $b_i \neq \theta_i$ is weakly dominated. There are two cases:

1. i bids $b_i > \theta_i$
 - (a) If there exists $b_j > b_i$ then $u_i(b_i, b_{-i}) = u_i(\theta_i, b_{-i})$
 - (b) If b_i is the highest bid, let b_j be the second highest. If $\theta_i < b_j < b_i$ then $u_i(b_i, b_{-i}) = \theta_i - b_j < 0 = u_i(\theta_i, b_{-i})$. If $b_j < \theta_i < b_i$ then again $u_i(b_i, b_{-i}) = u_i(\theta_i, b_{-i})$.
2. i bids $b_i < \theta_i$
 - (a) If b_i is the highest bid, then $u_i(b_i, b_{-i}) = u_i(\theta_i, b_{-i})$
 - (b) Some other b_j is the highest bid. Then if $b_i < b_j < \theta_i$, $u_i(b_i, b_{-i}) = 0$ while $u_i(\theta_i, b_{-i}) \geq 0$; in all other cases $u_i(b_i, b_{-i}) = u_i(\theta_i, b_{-i})$

Thus, indeed bidding the true value θ_i is a weakly dominant strategy.

Part 2. The efficient allocation rule $k^*(\cdot) = (k_1^*(\cdot), \dots, k_n^*(\cdot))$ (where $k_i \in [0, 1]$ is the probability that i gets the item) is such that for every profile of valuations $v = (\theta_1, \dots, \theta_n)$: $k^*(v) \in \arg \max_{k(v)} \{\sum_{i=1}^n \theta_i k_i(v)\}$. It is easy to see that the efficient allocation rule then prescribes giving the item to the bidder with the highest valuation θ_i : $k_i^*(v) = \mathbb{I}\{\theta_i > \theta_j \forall j\}$ (coupled with any tie-breaking rule).

From part 1 we know that the second-price auction has an equilibrium in (weakly) dominant strategies, in which all bidders bid truthfully. It is immediate that in this equilibrium, the bidder with the highest valuation wins the item – i.e., the SPA DS-implements the efficient allocation rule.