Dynamic Mechanism Design without Transfers

by

Young Wu

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Economics
University of Toronto

Abstract

Dynamic Mechanism Design without Transfers

Young Wu

Doctor of Philosophy

Graduate Department of Economics

University of Toronto

2018

This thesis consists of three chapters.

In chapter 1, titled "Design of Committee Search," I apply a mechanism design approach to committee search problems, such as hiring by a department or a couple's search for a house. A special class of simple dynamic decisions rules have agents submit in each period one of three votes: veto, approve, or recommend; the current option is adopted whenever no agent vetoes and at least one agent recommends. I show that every implementable payoff can be attained by randomizing among these simple rules. This result dramatically simplifies the design problem.

In chapter 2, titled "School Choice with Observable Characteristics," I study a school choice problem where students have observable characteristics that are correlated with their preferences. For example, one such characteristic may be the location of a student's home, which is correlated with preferences if students tend to prefer nearby schools. I consider mechanisms that are envy-free, efficient, and treat students with the same observable characteristics equally. I show that the welfare-maximizing mechanism in this class is a modified probabilistic serial mechanism with capacities. These capacities specify the maximum number of students with given characteristics that can be admitted into each school.

In chapter 3, titled "Mechanism Design for Stopping Problems with Two Actions," I analyse a class of dynamic mechanism design problems in which a single agent privately observes a time-varying state, chooses a stopping time, and upon stopping, chooses between two actions. The principal designs transfers that depend only on the time the agent stops and on the alternative the agent chooses. The analysis provides necessary and sufficient conditions for implementability in this environment. In particular, I show that any stopping rule in which the agent stops the first time the state falls outside of an interval in the state space can be implemented if and only if a pair of monotonicity conditions is satisfied.

Acknowledgements

I would like express my sincere gratitude to my advisors Dr. Marcin Peski, Dr. Colin Stewart, and Dr. Xianwen Shi for their support of my Ph.D study and related research, and for their help and guidance in the writing of this thesis.

Contents

Chapter 1

Design of Committee Search

1.1 Introduction

A hiring committee with two members receives job applications and conducts interviews until a position is filled. A decision is made right after each interview and is irreversible. Every committee member obtains private value from hiring a candidate. A common decision rule is the unanimity rule, according to which a candidate is hired whenever he is preferred by both members. Another possibility is a rule according to which each member submits a numerical score for a candidate and hiring occurs whenever the sum or average score is above a preset threshold.

A married couple looks for a house until they decide on purchasing one. The Canadian housing market is competitive, and a house is gone before a new one becomes available. The husband and the wife have different views regarding the ideal house, and their preferences with respect to its style and size may differ. In many households, decisions are not made collectively, and one member can establish dictatorship. In other cases, each member of the couple rates a house as ”ideal”, ”acceptable” or ”uninhabitable”, and they purchase the house if it is not uninhabitable for either member and ideal for at least one.

Two coauthors periodically gain access to new data sets. The authors have differing opinions about whether a certain data set could lead to interesting results. Data sets are costly to obtain and process, and they are likely to be used by other authors before a new one appears. A reasonable decision rule could be adopting a data set whenever at least one of the coauthors thinks the potential results are interesting. The threshold for a data set to be deemed interesting can vary between coauthors and and over time.

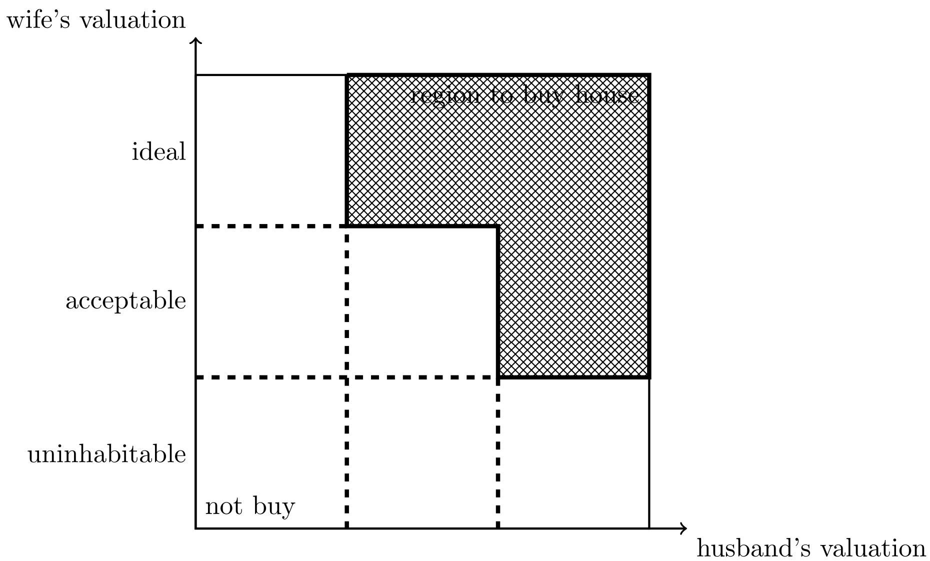

Figure 1.1: A possible decision rule in the couple's housing example

In all of these scenarios, two agents have to make a collective and irreversible decision regarding an option to adopt when facing options that are presented to them sequentially. Monetary transfers are not feasible. All the decision rules in these examples are plausible ones that are used widely in various collective search problems. In this paper, I focus on a class of simple voting rules, called ternary rules, that give each agent the power to veto, approve, or recommend an option, and the current option is adopted whenever no one vetoes and at least one of the agents recommends. It is similar to the "ideal-acceptable-uninhabitable" rule described in the house-searching example and in Figure 1.1. The shaded region in Figure 1.1 depicts the set of value pairs that lead to the purchase of the current house. The contribution of this paper is showing that every implementable decision rules is payoff-equivalent to some randomization among ternary rules, and as a result, a designer can restrict attention to the use of only these simple voting rules.

In particular, I consider a model in which two agents observe private values of, for example, hiring a candidate, in every period. The values are independent over time, but the values from different agents in the same period can be correlated. In each period, either the candidate is hired, in which case, every agent gets her private value, or the decision is delayed to the next period. The mechanism incentivizes the agents to report their values truthfully without using transfers, and it uses differing future decision rules to link incentives over time.

I assume a finite time horizon, and each agent obtains a fixed value from an outside option if no candidate is ever hired. In the last period, the only implementable decisions are,

- unanimity, where the candidate is hired when both agents have values higher than their outside options,

- dictatorship, where the candidate is hired as long as the dictator agent has value higher her outside option,
- reverse unanimity, where the candidate is not hired when both agents have values lower than their outside options, and
- constant, where the candidate is either always hired or never hired.

In the second to last period, a ternary rule specifies a lower threshold of value below which an agent vetoes and an upper threshold above which an agent recommends. It also needs to specify which rule to use in the last period after each pair of reported values. The following is an example of an implementable ternary rule.

- If one agent recommends and the other does not veto, the candidate is hired, and the search ends.
- If one agent recommends and the other vetoes, the candidate is not hired, and the agent who recommends becomes the dictator in the last period.
- If both agents choose to either approve or veto, the candidate is not hired, and the agents use the unanimity in the last period.

The reason for using the dictator rule is to give the agent who vetoes another agent's preferred candidate a lower expected future payoff as punishment to ensure that veto power will be exercised with discretion.

Many other decision rules are implementable in the periods before the last such as the one depicted in Figure 1.2. Given an implementable rule, the domain of each agent's value can be partitioned into three regions:

- the region where the candidate is not hired regardless of the other agent's value, called the veto region,
- the region where the candidate is hired as long as the other agent does not veto, called the recommendation region, and,
- the remaining region, called the approval region.

A candidate may or not may be hired in the region where both agents approve. Among the rules with the same partitions, the one that always rejects the candidate in the region where both agents approve

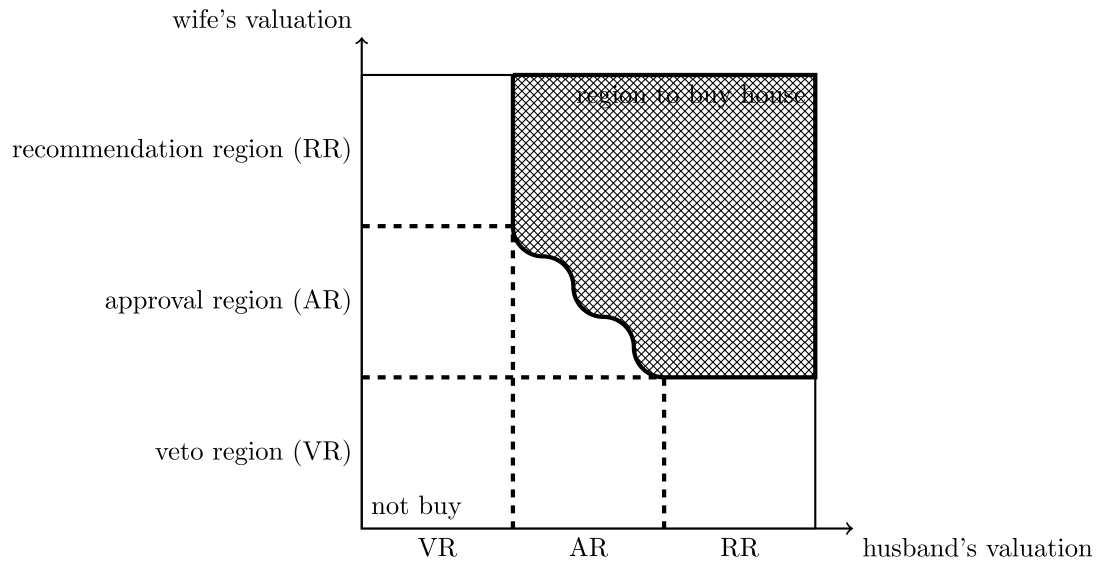

Figure 1.2: Another possible decision rule in the couple's housing example

is a ternary rule, and it yields the highest payoffs for both agents. This is the case because, intuitively, when considering a candidate whom both agents do not recommend, no agent will have a significant loss from continuing the search, and on average, they can both be better off if they continue.

For the same reason, the rule that always hires the candidate in the region where both agents approve is another ternary rule that yields lower payoffs for both agents. The dictator rule, according to which the candidate is hired whenever the dictator's value is above some threshold, is another special case of a ternary rule. With a carefully chosen threshold, dictatorship gives one agent higher payoff and the other agent lower payoff compared to the original rule. Therefore, any implementable decision rule is payoff-equivalent to the randomization between these ternary rules.

Technically, the model involves the implementation of an incentive compatible mechanism. Dynamic ex-post incentive compatibility is used as the definition of implementability so that the solution is robust to private communication between agents and robust to correlation between values and beliefs of the agents. I also restrict attention to mechanisms that are deterministic in every period, where either the candidate is hired for certain, or the decision is delayed. These assumptions make the model tractable. I make other assumptions for the purpose of presentation. For example, there are only two agents and a finite number of periods, and the private values are drawn from distributions with full support on some compact set and are independent over time. Those assumptions can be relaxed without significantly changing the main results.

This paper is closely related to the literature on dynamic mechanism design without transfers. Guo and Hörner [18] investigate the problem without transfers, and in which goods can be allocated in multiple periods. They characterize the optimal mechanism and provide a simple implementation in which the agents have virtual budgets. Lipnowski and Ramos [25] develop a similar model with similar results, in which the principal has less commitment power. In this model, the principal can decrease the veto threshold in the next period if she exercises her veto power. This is similar to the budget mechanism from Guo and Hörner [18] in the sense that the agent has a virtual budget of veto power, and if she vetoes excessively, she will gradually lose her power to veto.

Kovác, Krähmer, and Tatur [23] examine the stopping problem in which only a single good is allocated. In their model, the optimal mechanism is one in which the principal chooses different probabilities for assigning the good at different times to incentivize the agent's use of the stopping rule the principal prefers. I focus on allocation rules that are deterministic in each period, so instead of change probabilities, the principal changes the amount of control an agent has over the allocation in every period to incentivize an agent. The size of the region where an agent has more control over the allocation, for example, to veto or recommend it, is different in every period, and this is used in place of the probabilities in the deadline mechanism proposed by Kovác, Krähmer, and Tatur [23]. Both Guo and Hörner [18] and Kovác, Krähmer, and Tatur [23] focus on single-agent problems, whereas this paper focuses on problems with more than one agent.

Johnson [20] considers the problem with multiple agents, but the good is allocated to only one of the agents. He uses the promised utility approach where the principal promises future allocation rules that correspond to different expected payoffs to incentivize truth-telling. I use the same approach, but I consider the problem with a public good. In his model, the agents trade favors in a virtual market for decision rights, and as a result, they take turns obtaining their favorite private allocation. In this model, the agents also trade decision rights, but the goal is to wait for the public allocation that is preferable to every agent.

Moldovanu and Shi [27] study the committee search problem similar to mine. They solve the stationary voting rules that have a single threshold for each agent. Compte and Jehiel [12] and Albrecht, Anderson, and Vroman [3] develop a similar model and focus on majority voting rules. I consider all direct revelation mechanisms, and one result states that if decisions cannot be linked over time, then all implementable decision rules are the voting rules with a single threshold, as in Moldovanu and Shi [27]. Additional decision rules that are not voting rules can be implementable in this environment through the linking of decisions over time, but I show that every implementable rule is payoff-equivalent to some randomization among ternary voting rules, and this justifies the restriction to using only simple voting rules when solving committee search problems.

This paper is organized as follows. Section 1.2 introduces the model. Section 1.3 characterizes implementability. Section 1.4 states and explains the main result: every implementable payoff can be attained by randomizing among ternary rules. Section 1.5 concludes.

1.2 Model

In this section, I describe the agent's stopping problem and the principal's design problem. In particular, I define implementability as dynamic ex-post incentive compatibility, and explain why it is appropriate for this model.

1.2.1 Valuations

In this subsection, I describe the payoffs and the class of mechanisms.

I describe the model with two agents. A principal hires a new employee through a committee. In each period, a new candidate appears and each agent in the committee observes the value of hiring the candidate. The number of periods T is finite and there is no discounting. I assume that the values in period $t, v_t = (v_{t,1}, v_{t,2}) \in \mathcal{V} = [\underline{v}, \bar{v}]^2$ are independently distributed over time with continuous density f_t and full support $\mathcal{V}$, and that the means are normalized to 0 in each coordinate.

I use v^t to denote a history of values, or reports, from period 1 to period t,

$$\begin{aligned} v^t &= (v_1, v_2, ..., v_t) \\ &= ((v_{1,1}, v_{1,2}), (v_{2,1}, v_{2,2}), ..., (v_{t,1}, v_{t,2})). \end{aligned}$$

I assume independence over time but not between agents. Independence is not necessary for the results in this paper to hold, I make the assumption only for simpler presentation. In particular, the assumption simplifies the notations and the calculations in the examples. The value of hiring a candidate to each agent may be correlated.

Due to the revelation principle for deterministic mechanisms with ex-post constraints in Jarman, Meisner et al. [19], I restrict attention to direct mechanisms, in which the agents report only their own valuation in every period, which indicates whether to hire the candidate in each period. I define the mapping $q : \mathcal{V} \to [0, 1]$ as a stage mechanism after some history of reports, $v^{t-1} = (v_1, v_2, ..., v_{t-1})$. Here, $q(v_t|v^{t-1})$ is the conditional probability that the candidate is hired in period t if the agents report v_t after the history v^{t-1}, condition on the candidate not hired in periods before t. When the context is clear, I omit the history after which the stage mechanism is used and use the expression $q(v_t)$. The grand mechanism, Q, is the collection of stage mechanisms following every possible history,

$$Q : \bigcup_{t \leq T} \mathcal{V}^t \to [0, 1] .$$

The mechanism can depend on historical reports. As a result, the principal can choose different allocations in the future as rewards or punishments to incentivize truthful reports.

If no candidate is hired by the end of the last period, the agents get values from the outside option in period $T + 1, v_{T+1} = v^\star = (v_1^\star, v_2^\star)$, in the interior of $\mathcal{V}$. For each terminal history in the form $v^{T+1} = (v_1, v_2, ..., v_T, v^\star) \in \mathcal{V}^{T+1}$, the total probability of hiring the candidate should add up to 1,

$$\sum_{t=1}^{T+1} q\left(v_t\right) \prod_{s=1}^{t-1} \left(1 - q\left(v_s\right)\right) = 1,$$

where $q(v_{T+1}) = 1$ for $v_{T+1} = v^\star$.

Given a grand mechanism Q, let $w_{i,t}\left(Q; v^{t-1}\right)$ denote the ex-ante continuation value of agent i in period t after some history $v^{t-1} \in \mathcal{V}^{t-1}$, before $v_t, v_{t+1}, ..., v_T$ are realized, assuming both agents report truthfully.

$$w_{i,t}\left(Q; v^{t-1}\right) = \mathbb{E}\left[\sum_{s=t}^{T+1} q\left(v_s\right) \prod_{s'=1}^{s-1} \left(1 - q\left(v_{s'}\right)\right) v_{i,s}\right],$$

where $q(v_{T+1}) = 1$ and $v_{T+1} = v^\star$ with probability 1.

For a mechanism, Q, the *continuation value after history* v^{t-1} is a pair,

$$w_t\left(Q; v^{t-1}\right) = \left(w_{1,t}\left(Q; v^{t-1}\right), w_{2,t}\left(Q; v^{t-1}\right)\right).$$

In particular, the total payoff an agent, i, gets from a mechanism, Q, is $w_{i,1}(Q)$.

I also define two grand mechanisms, Q_1, Q_2, to be *payoff-equivalent* after history v^{t-1} if,

$$w_t\left(Q_1; v^{t-1}\right) = w_t\left(Q_2; v^{t-1}\right).$$

In particular, the payoff-equivalence between two dynamic mechanism Q_1 and Q_2 after history v^{t-1} is defined in terms of the ex-ante continuation payoff in period t before the vector of period-t valuations $(v_{1,t}, v_{2,t})$ is realized.

1.2.2 Implementability

In this subsection, I define dynamic ex-post implementability and outline the reasons it is used for this problem. Moreover, I define and restrict attention to a subset of random mechanisms I call quasi-deterministic mechanisms, where the mechanism is deterministic in each stage but between-period randomization is allowed.

Definition 1. A mechanism, Q, is (dynamic ex-post) incentive compatible, if in each period $t \in \{1, 2, ..., T\}$, after each history $v^{t-1} \in \mathcal{V}^{t-1}$, every $i \in \{1, 2\}$, every $v_{i,t} \in [\underline{v}, \bar{v}]$, and every $v_{-i,t} \in [\underline{v}, \bar{v}]$,

$$v_{i,t} \in \arg \max_{\hat{v}_{i,t} \in [\underline{v}, \bar{v}]} v_{i,t} q\left(\hat{v}_{i,t}, v_{-i,t}\right) + \left(1 - q\left(\hat{v}_{i,t}, v_{-i,t}\right)\right) w_{i,t+1}\left(Q; v^{t-1}, \hat{v}_{i,t}, v_{-i,t}\right).$$

A mechanism is dynamic ex-post incentive compatible (or the corresponding decision rule is dynamic ex-post implementable) if, in every period, it is optimal to report the true valuation given the other agent's report.

This definition is the same as the one from Noda [28], but without transfers and discounting. Noda [28] terms this type of implementability *within-period ex-post incentive compatibility,* or wp-EPIC. Bergemann and Välimäki [6], Parkes, Cavallo, Constantin, and Singh [29],and Athey and Segal [4] use a more complicated version of ex-post implementability since valuations and types in their models can be correlated over time.

Ex-post implementability is used because it is robust to private communication between agents, within-period correlation between valuations, and variation in agents' beliefs about each other's types. This definition of incentive compatibility also makes the model tractable. Bergemann and Morris [5] provide details regarding the use of ex-post incentive compatibility.

I omit individual rationality constraints, and agents are forced to participate.

I focus on mechanisms that are deterministic in every stage in which either the candidate is hired for certain or the decision is delayed to the next period.

Definition 2. A stage mechanism, q, is *deterministic* if $q(\cdot) \in \{0, 1\}$. A grand mechanism, Q, is *quasi-deterministic* if after every history $v^{t-1} \in \mathcal{V}^{t-1}$, only randomization among multiple deterministic stage mechanisms is used.

For example, after $(v_{1,t}, v_{2,t})$ are reported, the principal must either hire the candidate with probability 1 or delay hiring to the next period, but in the event that the principal chooses to delay, he can randomize among multiple stage mechanisms in period $t + 1$.

1.3 Implementability

In this section, I characterize the conditions for implementability. It is convenient to divide the analysis into a few steps. In the first step, I characterize implementable stage mechanisms in the last period. Next, I explain that the implementability in the dynamic setting can be reduced to a version of the static problem with appropriately chosen continuation values. Finally, I explain why the dynamic problem is substantially richer than a sequence of static ones.

1.3.1 Static Implementation

In this subsection, I consider the problem when $T = 1$. In this case, the continuation value is fixed at the outside option $(v_1^\star, v_2^\star)$. I describe a class of mechanisms called binary stage mechanisms that characterizes implementability. The same characterization applies to the last period when $T > 1$.

Definition 3. A stage mechanism, q, is binary, with outside option (a_1, a_2), if for each $i \in \{1, 2\}$, for every $v_{-i,t} \in [\underline{v}, \bar{v}]$, either $q(v_{i,t}, v_{-i,t})$ is constant (0 or 1) for every $v_{i,t} \in [\underline{v}, \bar{v}]$, or,

$$q(v_{i,t}, v_{-i,t}) = \begin{cases} 0 & \text{if } v_{i,t} < a_i \\ 1 & \text{if } v_{i,t} > a_i \end{cases}.$$

As shown in Figure 1.3, there are only six mechanisms that are binary for a fixed outside option. The shaded regions are the acceptance regions, $\{v_t : q(v_t) = 1\}$, where the candidate is hired. The remaining region is where the candidate is not hired and the outside option (a_1, a_2) is given to the agents.

Lemma 1. *For $T = 1$, if a quasi-deterministic mechanism, Q, is incentive compatible, then the (only) stage mechanism $q(\cdot|\emptyset)$ must be binary with outside option $(v_1^\star, v_2^\star)$.*

The intuition for Lemma 1 is as follows. Note that in each of the six binary stage mechanisms, whenever there is a threshold above which the candidate is hired and below which the candidate is not, the threshold value must be $v_1^\star$ for agent 1 and $v_2^\star$ for agent 2. This is because, in the last period, if the candidate is still not hired, the agents get $(v_1^\star, v_2^\star)$ in period $T + 1$. For an agent and a fixed report from the other agent, if the mechanism always hires the candidate or never hires the candidate, this agent will be indifferent between reporting any value, and as a result, she will not misreport. Otherwise, if the mechanism hires the candidate when she reports a value lower than the outside option, she will misreport a higher value and obtain the outside option instead; and if the mechanism does not hire the candidate when she reports a value higher than the outside option, she will misreport a lower value to

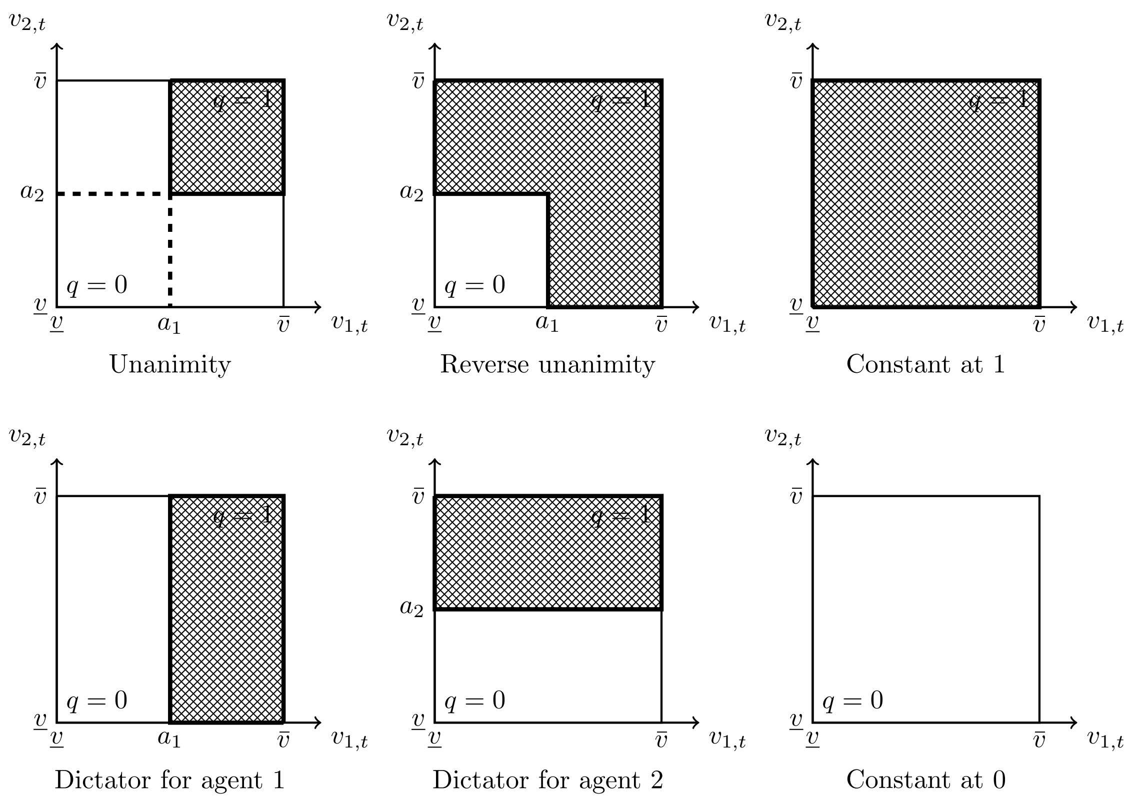

Figure 1.3: All six binary stage mechanisms

get the candidate hired. Therefore, the only incentive-compatible mechanism hires the candidate if and only if an agent reports a value higher than the outside option.

1.3.2 Dynamic Implementation

In this subsection, I show how the dynamic problem can be reduced to a static one but with more than one possible outside options. I then state a monotonicity condition as a characterization of all quasi-deterministic incentive-compatible mechanisms in this environment.

An incentive-compatible grand mechanism consists of a collection of incentive-compatible stage mechanisms, one after each history v^{t-1}, with continuation value pairs for all reports in the rejection region, $\{v_t : q(v_t) = 0\}$, where each continuation value pair corresponds to some sequence of incentive-compatible stage mechanisms in periods $t+1, t+2, ..., T$.

Given a stage mechanism $q\left(\cdot|v^{t-1}\right)$ and the report of every agent other than i, define the following threshold function,

$$R_i\left(v_{-i,t}\right) = \inf\left\{v_{i,t} : q\left(v_t\right) = 1\right\}, \tag{1.1}$$

with the convention that $\inf \emptyset = \infty$.

For deterministic stage mechanisms, there must thresholds (R_1, R_2) such that the candidate is hired if and only if an agent reports a value above the threshold. By incentive compatibility, the continuation value function must be constant whenever the value observed is below the threshold, since if it is not, an agent can misreport and obtain a different, and possibly higher, continuation value. This observation is stated as the following lemma.

Lemma 2. *A quasi-deterministic mechanism, Q, is incentive compatible if and only if after every history v^{t-1}, for each $i \in \{1,2\}$, and every $v_{-i,t} \in [\underline{v}, \bar{v}]$,*

$$q(v_t) = \begin{cases} 0 & \text{if } v_{i,t} < R_i(v_{-i,t}) \\ 1 & \text{if } v_{i,t} > R_i(v_{-i,t}) \end{cases},$$

and for every $v_t = (v_{i,t}, v_{-i,t})$ such that $q(v_t) = 0$,

$$w_{i,t+1}\left(Q; v^{t-1}, v_t\right) \text{ is independent of } v_{i,t}, \text{ and,}$$
$$w_{i,t+1}\left(Q; v^{t-1}, v_t\right) = R_i(v_{-i,t}) \text{ if } R_i(v_{-i,t}) \neq \infty.$$

The property is stated recursively: the payoff $w_{i,t+1}$ is the payoff from some sequence of stage mechanisms in periods $t+1, t+2, ...T$, and those stage mechanisms also satisfy the above conditions.

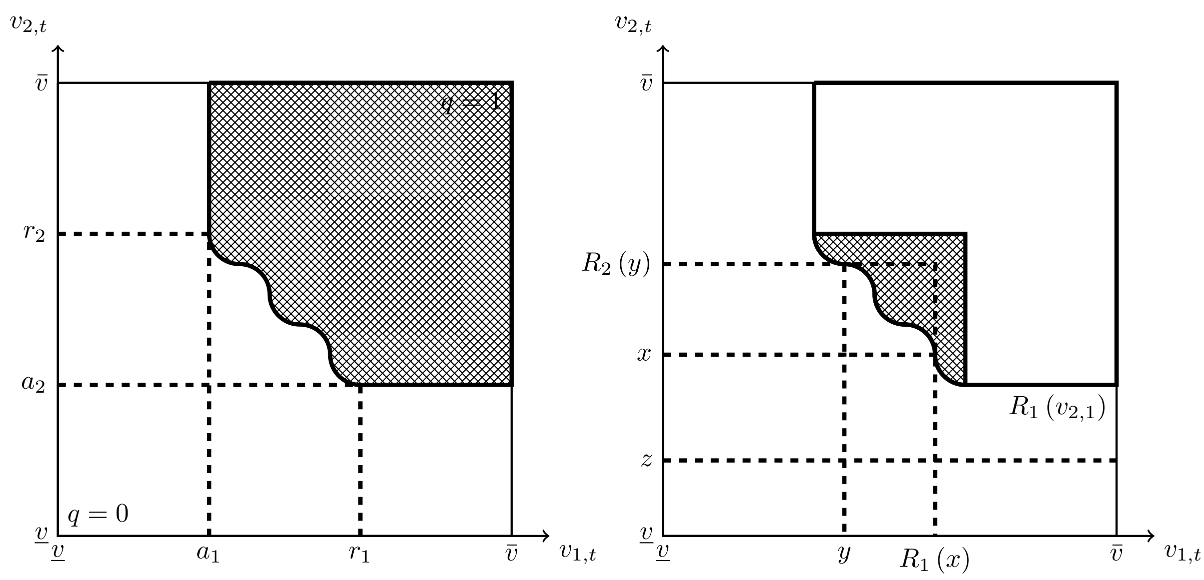

Figure 1.4: Example of a stage mechanism that is possibly incentive compatible

An example of a stage mechanism that satisfies the condition of Lemma 2 is depicted in the diagram on the left-hand side of Figure 1.4. According to the first part of Lemma 2, the acceptance region has the monotonicity property: if the candidate is accepted at some value pair v_t, he is also accepted at the value pairs that dominate $v_{t.}$

The continuation values for agent 1 along the line $v_{2,t} = z$ in the diagram on the right-hand side of Figure 1.4 are constant; for example,

$$w_{1,t+1}\left(v^{t-1},(v_{1,t},z)\right) = a_1 \; \forall \; v_{1,t} \in \left[\underline{v},\bar{v}\right].$$

If continuation values for agent 1 are not the same along this line, then the type of agent 1 who has a value that leads to a lower continuation value will misreport and get a different and higher continuation value.

The continuation values of agent 1 along $v_{2,t} = x$ in the diagram on the right-hand side of Figure 1.4 are equal to the lowest valuation for which the candidate is accepted $R_1(x)$, or,

$$w_{1,t+1}\left(v^{t-1},(v_{1,t},x)\right) = R_1(x) \;\; \forall \; v_{1,t} \text{ such that } q(v_{1,t},x) = 0.$$

Not only must the continuation value be constant for every $v_{1,t}$, it must be equal to $R_1(x)$, because if the continuation value is higher, say $R_1(x)+\delta$, then the type of agent 1 with value $R_1(x)+\dfrac{\delta}{2}$ will have an incentive to misreport a lower value and delay hiring to get a higher continuation value. Similarly, if the continuation value is lower, say $R_1(x)-\delta$, then an agent 1 observing $R_1(x)-\dfrac{\delta}{2}$ will have an incentive to misreport a higher value to get the candidate hired in the current period.

Similar requirements apply to agent 2 along the vertical line segments in the diagram. For example, the continuation value at the point (x,y) is fixed for both players at,

$$w_{t+1}\left(v^{t-1},(x,y)\right) = \left(R_1(y), R_2(x)\right).$$

1.3.3 Linking Decisions

The principal can link payoffs over time by using mechanisms that are history dependent. In this subsection, I demonstrate, with two examples, that such mechanisms can lead to Pareto improvements over history-independent ones.

As demonstrated in Figure 1.4, there are many incentive-compatible stage mechanisms in periods $1,2,...,T-1$ that look different from the ones in the last period, T. A key reason that more allocation rules are implementable is the linking of decisions: the principal can use different randomizations among stage mechanisms after different reports to punish or reward the agents. I call these mechanisms history-dependent mechanisms.

Definition 4. A grand mechanism is *history independent* if the distributions over the stage mechanisms after any two histories with the same length are the same, that is, if

$$q(\cdot|v^t) = q(\cdot|\tilde{v}^t) \text{ for every } v^t, \tilde{v}^t \in \mathcal{V}^t, \text{ and every } t \in \{1, 2, ..., T\}.$$

In general, if the principal is restricted to the use of history-independent mechanisms, then every stage mechanism in every period must be binary. This observation is stated in the following corollary to Lemma 1.

Corollary 1. *If a quasi-deterministic incentive-compatible mechanism is history independent, then the stage mechanisms chosen with strictly positive probabilities after every history must be binary.*

Proof. After every history v^t, only one continuation value is allowed in a history-independent mechanism, because there is only one sequence of stage mechanisms in periods $t+1, t+2, ..., T$. Then, every period is similar to the last period, with the exception that the outside option may differ from $(v_1^\star, v_2^\star)$. The set of stage mechanisms that is incentive compatible with a single outside option is the set of binary stage mechanisms for that outside option. □

The example below with $T = 2$ illustrates the difference between a history-dependent and a history-independent mechanism.

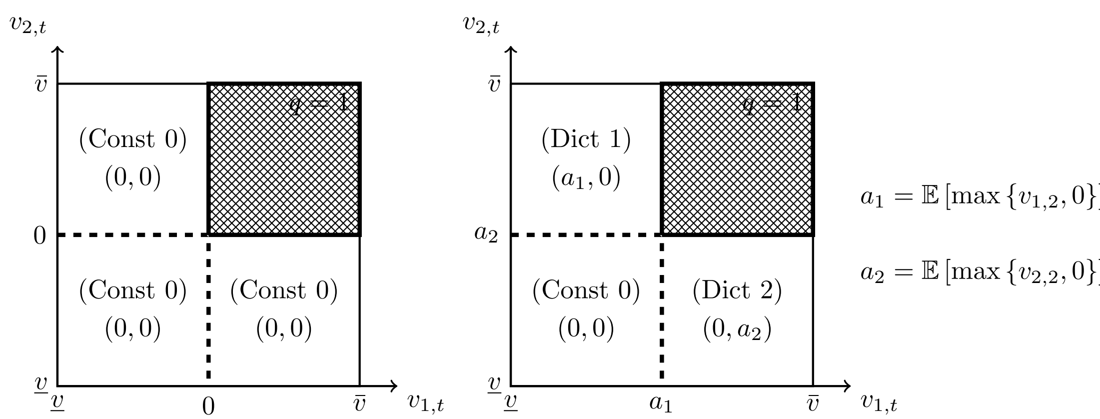

Figure 1.5: History independent (left) vs dependent (right) unanimity in period 1

Example 1. Suppose the value distributions are symmetric, the outside option is $v^\star = (0, 0)$, and the stage mechanism in period 1 involves symmetric unanimity rules. The two choices for the stage mechanisms in period 2 are

1. constant 0 (or unanimity or reverse unanimity) after all histories in which the candidate is not hired in period 1 or,

2. dictatorship by agent 1 after histories in which the candidate is not hired because agent 1 reports a value lower than a_1 but agent 2 reports a value higher than a_2; dictatorship by agent 2 after histories in which the candidate is not hired because agent 2 reports a value lower than a_2 but agent 1 reports a value higher than a_1; and constant 0 in period 2 after histories in which the candidate is not hired because both agents report values lower than their respective $a_{i.}$

These two mechanisms are depicted in the diagrams in Figure 1.5. Mechanism (1) is history independent and mechanism (2) is history dependent. Both diagrams depict the unanimity stage mechanism in period 1, and the name of the stage mechanism and its resulting continuation value pair from period 2 is written in each region in which that continuation mechanism is used.

Note that $a_1 = \mathbb{E}[\max\{v_{1,2}, 0\}]$ and $a_2 = \mathbb{E}[\max\{v_{2,2}, 0\}]$ because the property in Lemma 2 must be satisfied and the $(\mathbb{E}[\max\{v_{1,2}, 0\}], 0)$ and $(0, \mathbb{E}[\max\{v_{2,2}, 0\}])$ are the continuation value generated by the two dictator stage mechanisms in the second period.

In the next example, I show that there are rules that are not binary and give both agents higher payoffs than any binary rules.

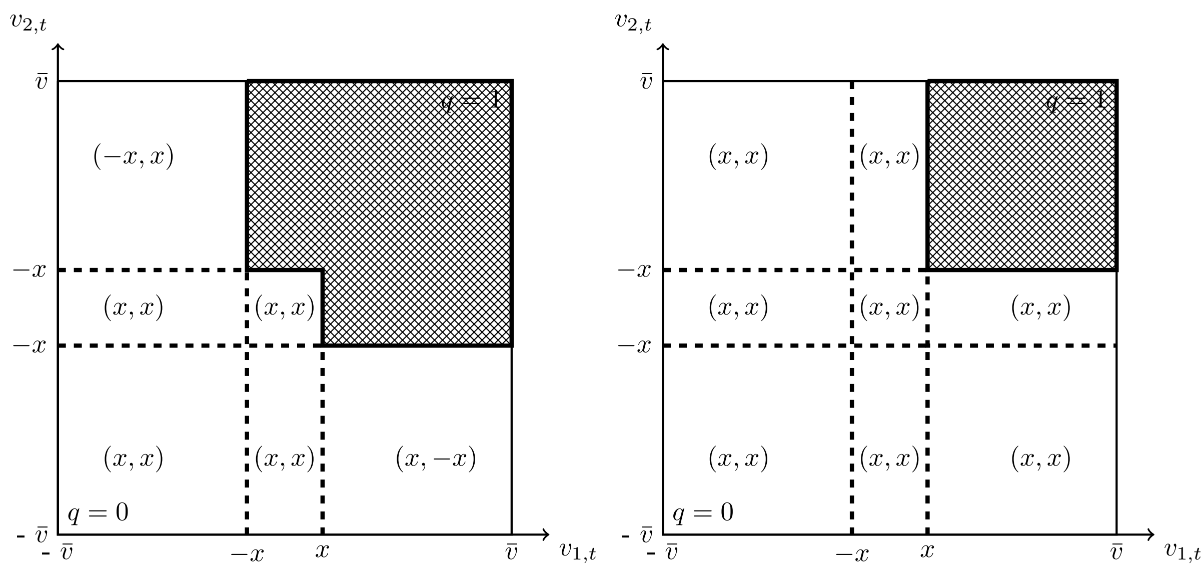

Figure 1.6: A mechanism, not binary, (left) that yields higher payoffs for both agents than the binary mechanism (right)

Example 2. Suppose the values in some period $t < T$ are independently piecewise uniformly distributed such that the probability in each of the nine rectangular regions in the diagram on the left-hand side of Figure 1.6 is $\frac{1}{9}$. I assume that in period $t+1$, the highest possible symmetric payoff pair that corresponds to some incentive-compatible mechanism is (x, x). For simplicity, let $\underline{v} = -\bar{v}$ and suppose $x < \frac{1}{7}\bar{v}$.

The difference in period t agent i payoff from the two stage mechanisms in Figure 1.6, after the same history v^{t-1}, is,

$$\begin{aligned}
& w_{i,t}\left(Q^{left}; v^{t-1}\right) - w_{i,t}\left(Q^{right}; v^{t-1}\right) \\
&= \frac{1}{9}\left(-x + \frac{x+\bar{v}}{2}\right) - \frac{1}{9}\left(2x + x\right) \\
&> \frac{1}{9}\left(-x + \frac{x+7x}{2}\right) - \frac{1}{9}\left(2x + x\right) \\
&= 0.
\end{aligned}$$

It can be shown that this stage mechanism yields higher payoff for both agents than do the other symmetric binary stage mechanisms, including the reverse unanimity and constant rules. In the next section, I show that this observation is not a coincidence and that all Pareto optimal stage mechanisms have the same shape as the one depicted in the diagram on the left-hand side of Figure 1.6.

1.4 Ternary Mechanisms

This section presents the main result of the paper. I start by defining ternary mechanisms, an important class of mechanisms with a very simple interpretation. The main result shows that any quasi-deterministic mechanism can be constructed by randomization among ternary stage mechanisms. The proof is divided into two steps. In the first subsection, I show that every Pareto optimal stage mechanism must be ternary. In the second subsection, I show that the rest of the boundary of the set of payoffs that can be generated by quasi-deterministic incentive-compatible mechanisms, that is not Pareto optimal, consists of only randomizations among binary mechanisms. The main result is stated and discussed in the last subsection.

Definition 5. A stage mechanism, q, is ternary if there are a_1, r_1, a_2, r_2, where $\underline{v} \leq a_i \leq r_i \leq \bar{v}$ for each $i \in \{1, 2\}$ such that,

$$q\left(v_t\right) = \begin{cases} 0 & \text{if } v_{i,t} < a_i \text{ for some } i \text{ or } a_i \leq v_{i,t} < r_i \text{ for all } i \\ 1 & \text{if } v_{i,t} > r_i \text{ for some } i \text{ and } v_{i,t} > a_i \text{ for all } i \end{cases}.$$

Figure 1.7 illustrates a typical ternary stage mechanism. The continuation values are placed in each rectangular region instead of the names of the continuation stage mechanisms from the next period. For general $T > 2$, the complete sequence of stage mechanisms after the current period is irrelevant and difficult to state explicitly; therefore, I write only the pair of continuation values, such as (r_1, r_2), in

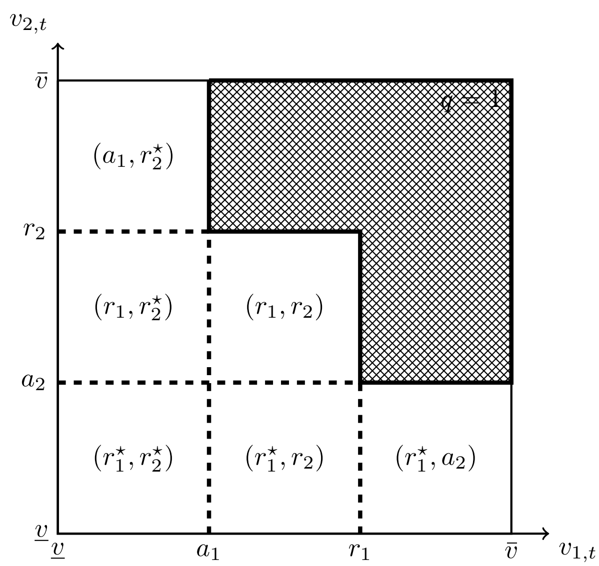

Figure 1.7: An example of a ternary stage mechanism

each rectangular region. This means that if the agents report v_t in this region, the candidate will not be hired and a sequence of stage mechanisms that results in payoff $w_t\left(Q; v^{t-1}\right) = (r_1, r_2)$ will be used in periods $t+1, t+2, ..., T$. As shown in Figure 1.7, the continuation values without a star are fixed due to Lemma 2, and the values with a star simply examples of possible continuation values.

Ternary stage mechanisms are voting mechanisms in which each agent can cast one of three votes: veto, approve, or recommend. Every agent can veto the candidate, and in order to hire the candidate, the principal needs at least one person to recommend the candidate. In particular, this voting rule does not hire the candidate if both agents approve and neither recommends. Each agent has three intervals separated by a_i and r_i, the smallest (lower than a_i) where the agent can veto the candidate, the largest (higher than r_i) where the agent recommends and the candidate is hired as long as the other agent does not veto, and the one in the middle where the candidate is hired only when the other agent recommends.

1.4.1 Pareto Optimal Mechanisms

In this subsection, I define Pareto optimal mechanisms as the ones that are restricted Pareto optimal within the set of quasi-deterministic incentive-compatible mechanisms. I briefly explain why if a stage mechanism is not ternary after some history, then it is Pareto dominated by one that is ternary after the same history.

Definition 6. For incentive-compatible mechanisms Q and $\tilde{Q}$, the mechanism Q *Pareto dominates* $\tilde{Q}$

after history v^{t-1} if

$$w_{i,t}\left(Q;v^{t-1}\right) \geq w_{i,t}\left(\tilde{Q};v^{t-1}\right) \text{ for each } i \in \{1,2\},$$

with strict inequality for at least one agent.

An incentive-compatible mechanism, Q, is *Pareto optimal after history* v^{t-1}, if it is not Pareto dominated by any other incentive compatible mechanism, $\tilde{Q}$, after the same history v^{t-1}.

Lemma 3. *If a stage mechanism is Pareto optimal after some history, then it is payoff-equivalent to a randomization among ternary stage mechanisms after the same history.*

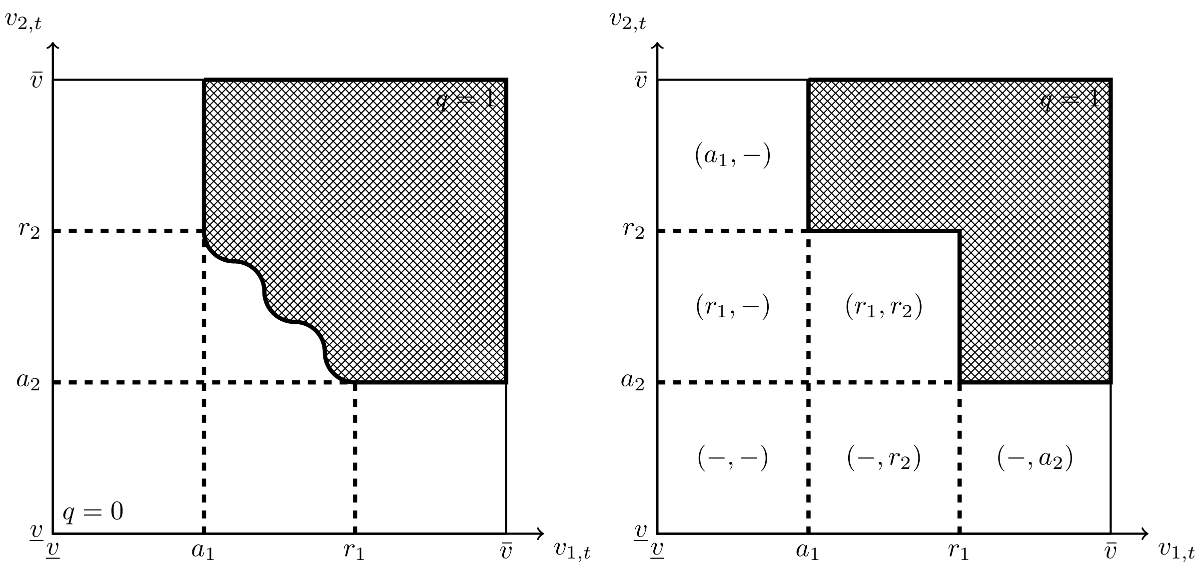

Figure 1.8: An arbitrary mechanism (left) and an ternary mechanism that Pareto dominates it (right)

For an arbitrary incentive-compatible stage mechanism like the one in the diagram on the left-hand side of Figure 1.8, I define a_i and r_i as follows:

$$a_i = \sup_{v_{i,t}} \left\{v_{i,t} : q\left(v_{i,t}, v_{-i,t}\right) = 0 \text{ for all } v_{-i,t} \in \mathcal{V}_{-i}\right\}, \tag{1.2}$$

If agent i observes a value lower than a_i, the candidate will not be hired, regardless of what value the other agent reports.

$$r_i = \lim_{\varepsilon \to 0^+} R_i\left(a_{-i} + \varepsilon\right), \tag{1.3}$$

If agent i observes a value higher than r_i, the candidate will be hired as long as the other agent reports a value higher than a_{-i}. Recall that R_i is the threshold function defined in Equation 1.1.

In the case in which the acceptance region is closed,

$$r_i = R_i\left(a_{-i}\right) < \infty. \tag{1.4}$$

I call a_i the approval threshold and r_i the recommendation threshold for an agent i. When agent i observes a value lower than a_i, the candidate will never be hired for any value the other agent reports, meaning agent i vetoes hiring. When agent i observes a value higher than r_i, the candidate will be hired as long as the other agent does not veto, meaning agent i recommends hiring. When both agents approve but none recommends, the stage mechanism can choose an arbitrary hiring rule such as the one in the left-hand-side diagram in Figure 1.4.

I compare the original stage mechanism in the diagram on the left-hand side of Figure 1.8 with one in which the candidate is never hired in the region $[a_1, r_1] \times [a_2, r_2]$ on the right-hand side of Figure 1.8. The latter stage mechanism is ternary. In the ternary mechanism, when $v_{2,t}$ is between a_2 and r_2, agent 1 gets constant continuation value a_1 if the candidate is not hired, whereas in the original mechanism, agent 1 gets continuation values that are lower than or equal to a_1 whenever she has a value lower than a_1. When $v_{2,t}$ is lower than a_2 or higher than r_2, the ternary mechanism yields the same expected payoff for agent 1 by construction.

1.4.2 Non-Pareto Optimal Mechanisms

The boundary of the set of continuation values that correspond to some incentive-compatible mechanism consists of Pareto optimal mechanisms and the ones that are the worst for one agent fixing the payoff of the other agent. In this subsection, I explain why the part of the boundary that is not Pareto optimal is made up of ternary mechanisms and conclude that every incentive-compatible mechanism is payoff-equivalent to some randomization among ternary mechanisms.

For an incentive-compatible mechanism, Q, it is *on the non-Pareto optimal boundary* after history v^{t-1}, if for some $i \in \{1,2\}$, there does not exist another incentive-compatible mechanism, $\tilde{Q}$, with the property that

$$\begin{aligned} w_{i,t}\left(\tilde{Q}; v^{t-1}\right) &= w_{i,t}\left(Q; v^{t-1}\right) \text{ and,} \\ w_{-i,t}\left(\tilde{Q}; v^{t-1}\right) &< w_{-i,t}\left(Q; v^{t-1}\right). \end{aligned}$$

The shaded region in Figure 1.9 represents a set of continuation values that are achievable by some incentive-compatible mechanisms. The Pareto optimal boundary is highlighted, and the remaining

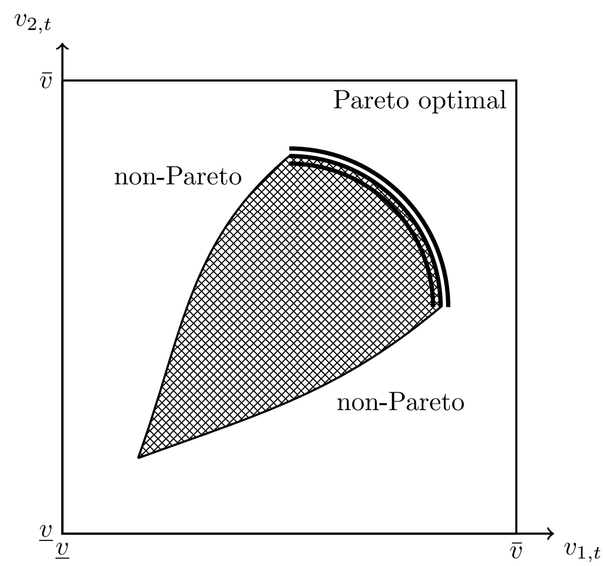

Figure 1.9: Boundary of a set of continuation values that corresponds to some incentive-compatible mechanism

boundary is the non-Pareto optimal boundary.

Lemma 4. *If the grand mechanism is on the non-Pareto optimal boundary after some history, then the stage mechanism after that history is payoff-equivalent to a randomization among binary stage mechanisms after the same history.*

I show that for an arbitrary stage mechanism like the one shown in Figure 1.8, a binary mechanism such as the one depicted in the left-hand-side diagram in Figure 1.10, with carefully constructed continuation value pairs in each rectangular region, results in a larger payoff for one agent but a smaller one for the other agent. Intuitively, the stage mechanism in the diagram is close to the dictator mechanism for agent 1, so it leads to a payoff that is higher for agent 1 and lower for agent 2. A ternary mechanism depicted in the left-hand side diagram in Figure 1.10 results in a smaller payoff for both agents. As explained at the end of the previous subsection, when both agents approve, never hiring the candidate is better for both agents than sometimes hire and sometimes not. For the same reason, always hiring is worse for both agents.

Theorem 1. *For every quasi-deterministic incentive-compatible mechanism, Q, there is a mechanism $\tilde{Q}$ that satisfies,*

1. *Every stage mechanism of $\tilde{Q}$ is ternary, and*

2. *Q is payoff-equivalent to $\tilde{Q}$.*

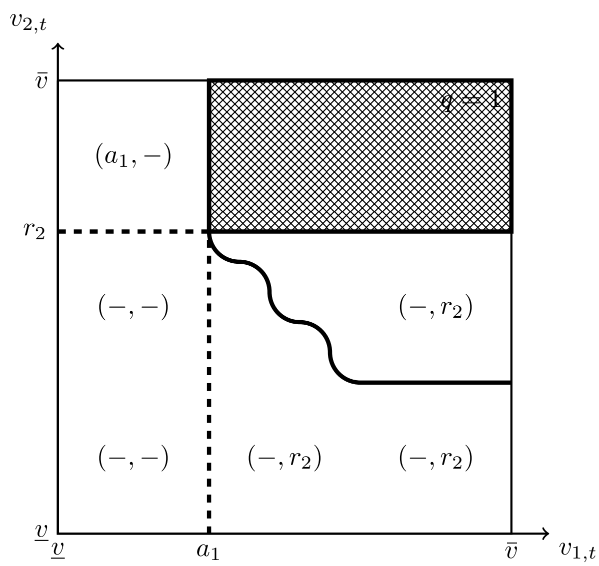

Figure 1.10: A binary mechanism better for 1 and worse for 2 (left) and a binary mechanism worse for both (right)

The result follows directly from Lemma 3 and Lemma 4, since the best and worst payoff for an agent are randomizations between ternary stage mechanisms, and this result applies to every agent and the stage mechanism after every history. Theorem 1 implies that a principal can use only ternary mechanisms to obtain every implementable payoff for the agents.

1.5 Discussion

In this section, I discuss possible extensions of the model.

The assumptions on the distributions over time are imposed for simpler presentation in the main text and the proofs, and they can easily be relaxed without significantly changing the proofs. In particular, the independence assumption is never used in the proofs due to the simplified definition of implementability used in the paper. However, a more complex definition of dynamic ex-post incentive compatibility should be used in the case with correlated values over time.

The definition of binary stage mechanisms can easily be extended to problems with more than two agents. These binary stage mechanisms remain the only incentive-compatible ones in the static case and in the last period for problems with $N > 2$ agents. The proofs of Lemma 1 and Corollary 1 applies directly to the general N agent problems. The definition of ternary stage mechanisms extend to problems with more than two agents. Lemma 3 and Lemma 4 apply also to the general agents problems, but since the proof requires additional notations that are complicated and do not provide additional insights, it is therefore not presented. However, the main result, Theorem 1, does not follow from Lemma 3 and Lemma 4, as in the two-agent case. This is because the mechanisms from Lemma 4 do not cover the entire non-Pareto boundary in N dimensions. Lemma 4 implies only that fixing the payoff of *one* agent makes it possible to decrease the payoffs of *all other* agents at the same time. I conjecture that Theorem 1 holds for N agent problems but that a different approach to the proof of Lemma 4 is required.

Another restrictive assumption is the requirement to use only quasi-deterministic mechanisms. A decomposability result similar to that provided by Pycia and Ünver [36] is required to show that any random mechanism can first be decomposed into multiple quasi-deterministic mechanisms and then constructed through further randomization among ternary mechanisms. Decomposability is not obvious in this model for even the simplest random mechanisms; therefore, I restrict attention to quasi-deterministic mechanisms without full justification.

It is possible to characterize the Pareto frontier in very simple examples. For example, if the valuations are independently uniformly distributed on [-1, 1], the Pareto frontier in the first period can be described by the line segment connecting the payoffs from the dictator mechanisms where one agent is the always the dictator in all periods. For two period problems, it is also possible to find the Pareto frontier numerically, but it is very difficult in general to provide a characterization of the Pareto frontier. As a result, it is also difficult to find the optimal dynamic mechanism given an objective function. They are interesting questions that I am unable to answer at the moment and I will leave them to potential future research.

1.6 Proofs (for the Lemmas in Section 3)

In this section, I start by stating the monotonicity condition that is necessary but not sufficient for implementability. Then I provide the proof of Lemma 2, and I use the result to prove Lemma 1.

Lemma 5. *If a grand mechanism, Q, is incentive compatible, then after every history $v^{t-1}, q\left(v_t|v^{t-1}\right)$ is weakly increasing in $v_{i,t}$ for each agent i.*

Proof of Lemma 5: Let v^t and $\tilde{v}^t$ be two histories that are the same except for the value for agent i in period $t, v_{i,t}$ and $\tilde{v}_{i,t}$, respectively, such that $v_{i,t} > \tilde{v}_{i,t}.$

Ex post implementability implies, for a stage mechanism $q\left(\cdot\right) = q_{\cdot|v^{t-1}}$ of Q,

$$q\left(v_t\right) v_{i,t} + \left(1 - q\left(v_t\right)\right) w_{i,t}\left(Q; v^t\right) \geq q\left(\tilde{v}_t\right) v_{i,t} + \left(1 - q\left(\tilde{v}_t\right)\right) w_{i,t}\left(Q; \tilde{v}^t\right) \text{ and },$$
$$q\left(v_t\right) \tilde{v}_{i,t} + \left(1 - q\left(v_t\right)\right) w_{i,t}\left(Q; v^t\right) \leq q\left(\tilde{v}_t\right) \tilde{v}_{i,t} + \left(1 - q\left(\tilde{v}_t\right)\right) w_{i,t}\left(Q; \tilde{v}^t\right).$$

Taking the difference between the above inequalities gives the result.

$$\begin{aligned}(q(v_t) - q(\tilde{v}_t)) \cdot (v_{i,t} - \tilde{v}_{i,t}) &\geq 0 \\ q(v_t) - q(\tilde{v}_t) &\geq 0\end{aligned}$$

Since v^t and $\tilde{v}^t$ only differ in the component $v_{i,t} > \tilde{v}_{i,t}, q$ is monotonic in the component $v_{i,t}$ for every i and every t. □

Proof of Lemma 2: Fix a grand mechanism Q and its stage mechanism after the history $v^{t-1}, q(\cdot) = q\left(\cdot|v^{t-1}\right)$, recall the definition of the threshold function in Equation 1.1, $R_i(v_{-i,t}) = \inf\{v_{i,t} : q(v_t) = 1\}$, with the convention that $\inf \emptyset = \infty$.

I first prove that the conditions (monotonicity on q and threshold condition on w_{t+1}) are necessary. Assume Q is a quasi-deterministic incentive compatible mechanism.

By Lemma 5, since q is monotonic in $v_{i,t}$, and,

$$q(v_{i,t}, v_{-i,t}) = \begin{cases} 0 & \text{if } v_{i,t} < R_i(v_{-i,t}) \\ 1 & \text{if } v_{i,t} > R_i(v_{-i,t}) \end{cases}.$$

The continuation value for agent i for fixed $v_{-i,t}, w_{i,t+1}\left(Q; v^{t-1}, (v_{i,t}, v_{-i,t})\right)$, must be constant for all values $v_{i,t}$ in the rejection region $\{v_t : q(v_t) = 0\}$. If not, the agent could always find it optimal to report the valuation with the largest possible continuation value. From now on, fix $v_{-i,t}$ and let $w_{i,t+1}\left(Q; v^{t-1}, (v_{i,t}, v_{-i,t})\right) = c$ be the constant continuation value.

For an agent observing $v_{i,t} < R(v_{-i,t})$, incentive compatibility implies reporting $v_{i,t}$ is preferred to reporting another $\hat{v}_{i,t} > R_i(v_{-i,t})$ to get the allocation $q = 1$, meaning,

$$c \geq v_{i,t}.$$

This is true for every $v_{i,t} < R(v_{-i,t})$,

$$c \geq R_i(v_{-i,t}).$$

For an agent observing $v_{i,t} > R_i(v_{-i,t})$, incentive compatibility implies reporting $v_{i,t}$ is preferred to reporting another $\hat{v}_{i,t} < R_i(v_{-i,t})$ to get the allocation $q = 0$, meaning,

$$c \leq v_{i,t}.$$

This is true for every $v_{i,t} > R(v_{-i,t})$,

$$c \leq R_i(v_{-i,t}).$$

Therefore,

$$c = R_i(v_{-i,t}),$$

and,

$$w_{i,t+1}\left(Q; v^{t-1}, (v_{i,t}, v_{-i,t})\right) = R_i(v_{-i,t}) \text{ for each } v_{i,t} \text{ such that } q(v_{i,t}, v_{-i,t}) = 0.$$

Now I prove that the monotonicity and threshold conditions are sufficient. Fix $v_{-i,t}$ and assume q is a stage mechanism that satisfy these conditions.

For an agent i with $v_{i,t} < R(v_{-i,t})$, reporting $\hat{v}_{i,t}$ results in payoff,

$$w_{i,t}\left(Q; v^{t-1}\right) = \begin{cases} R_i(v_{-i,t}) & \text{if } \hat{v}_{i,t} = v_{i,t} \\ R_i(v_{-i,t}) & \text{if } \hat{v}_{i,t} \neq v_{i,t}, \hat{v}_{i,t} \leq R(v_{-i,t}) \\ v_{i,t} & \text{if } \hat{v}_{i,t} > R(v_{-i,t}) \end{cases}.$$

Since $v_{i,t} < R_i(v_{-i,t})$, it is optimal to report truthfully.

For an agent i with $v_{i,t} > R(v_{-i,t})$, reporting $\hat{v}_{i,t}$ results in payoff,

$$w_{i,t}\left(Q; v^{t-1}\right) = \begin{cases} v_{i,t} & \text{if } \hat{v}_{i,t} = v_{i,t} \\ v_{i,t} & \text{if } \hat{v}_{i,t} \neq v_{i,t}, \hat{v}_{i,t} > R(v_{-i,t}) \\ R_i(v_{-i,t}) & \text{if } \hat{v}_{i,t} \leq R(v_{-i,t}) \end{cases}.$$

Since $v_{i,t} > R_i(v_{-i,t})$, it is optimal to report truthfully.

Therefore, truthful reports are optimal, q is incentive compatible. $\square$

Proof of Lemma 1: Fix a grand mechanism, Q. Since for any history $v^T \in \mathcal{V}^T$,

$$w_{i,T+1}\left(Q; v^T\right) = v_i^\star,$$

by Lemma 2, for any $v_{-i,T}$, either $q\left(v_{i,T}, v_{-i,T} | v^{T-1}\right)$ is constant in $v_{i,T}$ or,

$$R_i\left(v_{-i,T}\right) = v_i^\star.$$

By Definition 3, these stage mechanisms are binary with outside option $v_i^\star$ for agent i. □

1.7 Proofs (for the Lemmas and Propositions in Section 4)

In this section, I start by defining some shorthand notations for the proofs in this section and some preliminary observations that simplify the shapes and continuation values of an incentive compatible mechanism. Then, I prove Lemma 3 and Lemma 4. Theorem 1 follows directly from Lemma 3 and Lemma 4.

From now on, I am going to fix the grand quasi-deterministic incentive compatible mechanism Q, period t and history v^{t-1}. To simplify the subsequent notation, I write,

$$q\left(v_t\right) = q\left(v_t | v^{t-1}\right).$$

I also write,

$$w_{t+1}\left(v_t\right) = w_{t+1}\left(Q; v^{t-1}, v_t\right),$$

and,

$$w_t = w_t\left(Q; v^{t-1}\right).$$

I define the following constants and sets, I assume the acceptance region $\{v_t : q\left(v_t\right) = 1\}$ is closed for these definitions.

1. Recall from Equation 1.2 that the *approval threshold* , a_i, is the threshold below which the agent i has veto power: if $v_{i,t} < a_i$, the candidate is vetoed by i and will never be hired for any value of

$v_{-i,t}$,

$$a_i = \sup_{v_{i,t}} \left\{ v_{i,t} : q\left(v_{i,t}, v_{-i,t}\right) = 0 \text{ for all } v_{-i,t} \in [\underline{v}, \bar{v}] \right\}.$$

2. The *non-threshold region* for agent i, is the set of $v_{-i,t}$ such that the candidate is never hired for any value of $v_{i,t}$,

$$\left\{ v_{-i,t} : q\left(\tilde{v}_{i,t}, v_{-i,t}\right) = 0 \text{ for all } \tilde{v}_{i,t} \in [\underline{v}, \bar{v}] \right\} = [\underline{v}, a_{-i}).$$

The *average non-threshold continuation value* , c_i, is the expected continuation value within the non-threshold region for agent i,

$$c_i = \mathbb{E}\left[w_{i,t+1}\left(v_{i,t}, v_{-i,t}\right) | v_{-i,t} < a_i \right].$$

3. The definition for the *recommendation threshold* , r_i, is simplified from Equation 1.4 due to the assumption that the acceptance region is closed,

$$r_i = R_i\left(a_{-i}\right).$$

One useful relation due to the definition of these constants is,

$$a_i \leq R_{v_{-i,t}} \leq r_i \text{ for all } v_{-i,t} \geq a_{-i}. \tag{1.5}$$

Then, I state an corollary to Lemma 2 that makes computation easier in all the following proofs.

Corollary 2. *Fix a history v^{t-1} and a stage mechanism of a quasi-deterministic incentive compatible Q after this history, $q\left(\cdot\right) = q\left(\cdot | v^{t-1}\right)$,*

$$\mathbb{E}\left[\max\left\{ v_{i,t}, R_i\left(v_{-i,t}\right) \right\} | v_{-i,t} \geq a_{-i} \right] \cdot \mathbb{P}\left\{ v_{-i,t} \geq a_{-i} \right\} + c_i \cdot \mathbb{P}\left\{ v_{-i,t} < a_{-i} \right\}.$$

Proof of Corollary 2: It follows directly from Lemma 2. □

In the course of this section, I will construct alternative stage mechanisms with new continuation values. I say *a continuation value is incentive compatible* if they can be implemented in an incentive compatible quasi-deterministic grand mechanisms.

As an example, I show the following Compactness Lemma that ensures that the acceptance region is closed.

Lemma 6. *There exists an incentive compatible stage mechanism $\tilde{q}$, with incentive compatible continuation values, that is payoff equivalent to q, such that,*

$$\{v_t : \tilde{q}(v_t) = 1\} \text{ is closed}.$$

Recall that payoff equivalence in this section means $\tilde{w}_t = w_t$ where $\tilde{w}_t$ is the continuation value of mechanism that has $\tilde{q}$ after history v^{t-1} in place of q.

Proof. Define $\tilde{q}$ such that the set $\{v_t : \tilde{q}(v_t) = 1\}$ is the closure of the set $\{v_t : q(v_t) = 1\}$. Payoff equivalence is due to continuity and full support property of the value distributions. □

The only issue is how to define the continuation payoffs for player i when the other player says a_i. But use the continuity and just take the limit of continuation values when $v_{-i} > a_i$.

Also, I show the following Flattening Lemma that ensures that continuation values in the non-threshold region is constant.

Lemma 7. *There exists an incentive compatible stage mechanism $\tilde{q}$, with incentive compatible continuation values, that is payoff equivalent to q, such that,*

$$\tilde{w}_{t+1}(v_{i,t}, v_{-i,t}) = c_i \text{ for all } v_{-i,t} < a_{-i}.$$

Proof of Lemma 7: Let $\tilde{q}$ be the stage mechanism from replacing the continuation value of agent i in q in the non-threshold region $v_{-i,t} < a_{-i}$ by c_i.

I divide the proof into three main parts.

1. $\tilde{q}$ is incentive compatible.

2. $\tilde{q}$ payoff equivalent to q.

Part (1) There are two types of new continuation values that is different from the ones used in q.

1. $(c_i, R_{-i}(v_{i,t}))$ from the region $v_{i,t} \geq a_i$ and $v_{-i,t} < a_{-i}$,

2. (c_i, c_{-i}) from the region $v_{i,t} < a_i$ and $v_{-i,t} < a_{-i}$.

Fix $\tilde{v}_{i,t} \geq a_i, R_{-i}(\tilde{v}_{i,t})$ is finite due to the definition of a_i.

$$\begin{aligned}(c_i, R_{-i}(\tilde{v}_{i,t})) &= (\mathbb{E}[w_{i,t+1}(v_{i,t}, v_{-i,t})\,|v_{-i,t} < a_{-i}], R_{-i}(\tilde{v}_{i,t}))\\ &= (\mathbb{E}[\mathbb{E}[w_{i,t+1}(v_{i,t}, v_{-i,t})\,|v_{-i,t} < a_{-i}], w_{-i,t+1}(\tilde{v}_{i,t}, v_{-i,t}))\\ &= \mathbb{E}[w_{t+1}(\tilde{v}_{i,t}, v_{-i,t})\,|v_{-i,t} < a_{-i}].\end{aligned}$$

The second equality is due to Lemma 2 which states that $w_{i,t+1}(v_{i,t}, v_{-i,t})$ is constant in $v_{i,t}$ in the non-threshold region $[\underline{v}, a_{-i})$.

Similarly,

$$\begin{aligned}(c_i, c_{-i}) &= (\mathbb{E}[w_{i,t+1}(v_{i,t}, v_{-i,t})\,|v_{-i,t} < a_{-i}], \mathbb{E}[w_{-i,t+1}(v_{i,t}, v_{-i,t})\,|v_{i,t} < a_{-i}])\\ &= \mathbb{E}[w_{t+1}(v_{i,t}, v_{-i,t})\,|v_{i,t} < a_{-i}, v_{-i,t} < a_{-i}].\end{aligned}$$

The second equality is due to Lemma 2 which states that $w_{i,t+1}(v_{i,t}, v_{-i,t})$ is constant in $v_{i,t}$ in the non-threshold region $[\underline{v}, a_{-i})$.

Therefore, $(c_i, R_{-i}(v_{i,t}))$ and (c_i, c_{-i}) are the expected value other continuation value pairs so they are incentive compatible too. Since $\tilde{q}$ is the same as q, so incentive compatible of $\tilde{q}$ follows from the incentive compatibility of q.

Part (2) Note that,

$$\begin{aligned}\tilde{c}_i &= \mathbb{E}[\tilde{w}_{i,t+1}(v_{i,t}, v_{-i,t})\,|v_{-i,t} < a_{-i}]\\ &= \mathbb{E}[c_i|v_{-i,t} < a_{-i}]\\ &= c_i.\end{aligned}$$

I check $\tilde{q}$ is payoff equivalent to q using Corollary 2,

$$\begin{aligned}\tilde{w}_{i,t} &= \mathbb{E}\left[\max\left\{v_{i,t}, \tilde{R}_i(v_{-i,t})\right\}|v_{-i,t} \geq a_{-i}\right] \cdot \mathbb{P}\{v_{-i,t} \geq a_{-i}\} + \tilde{c}_i \cdot \mathbb{P}\{v_{-i,t} < a_{-i}\}\\ &= \mathbb{E}[\max\{v_{i,t}, R_i(v_{-i,t})\}\,|v_{-i,t} \geq a_{-i}] \cdot \mathbb{P}\{v_{-i,t}v_{-i,t} \geq a_{-i}\} + c_i \cdot \mathbb{P}\{v_{-i,t} < a_{-i}\}\\ &= w_{i,t}.\end{aligned}$$

$\square$

Lemma 8. *Ternary stage mechanisms with continuation values described in diagram on the right of*

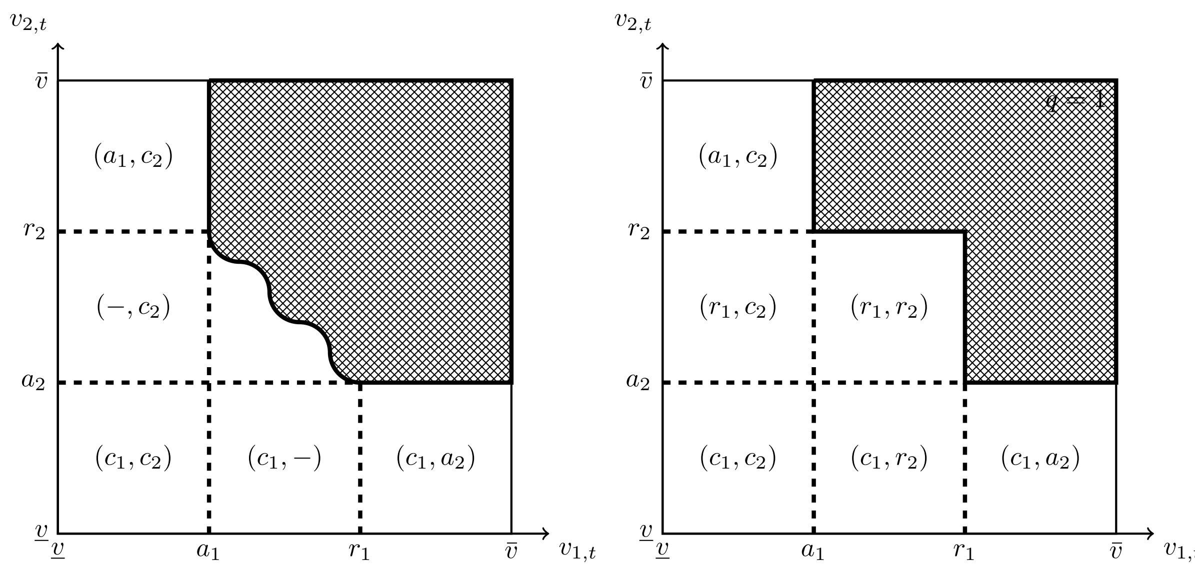

Figure 1.11: Mechanism $\tilde{q}$

Figure 1.11 are incentive compatible as long as the continuation values are incentive compatible.

Proof. It follows directly from Lemma 2. □

In the following proofs, by Lemma 6 and Lemma 7, without loss of generality, assume the acceptance region of q is closed and continuation value for each agent i in its non-threshold region $v_{-i,t} \in [\underline{v}, a_{-i})$ is constant at c_i. Let the continuation values and the shape of the acceptance region of q be ones as in .

Proof of Lemma 3: Consider a ternary stage mechanism, $\tilde{q}$, with the thresholds a_i, r_i with continuation values specified in .

I divide the proof into three main parts.

1. $\tilde{q}$ is incentive compatible.

2. $\tilde{q}$ Pareto dominates q.

Part (1) I only need to show that the continuation values are incentive compatible, the rest follows from Lemma 8. There are three types of new continuation values that is different from the ones used in q.

1. (r_i, r_{-i}) in the region $a_j \leq v_{j,t} < r_j$ for both $j \in \{1, 2\}$,

2. (r_i, c_{-i}) in the region $v_{i,t} < a_i$ and $a_{-i} \leq v_{-i,t} < r_{-i}$.

$$
\begin{aligned}
(r_i, r_{-i}) &= (R_i(a_{-i}), R_{-i}(a_i)) \\
&= w_{t+1}(a_i, a_{-i}),
\end{aligned}
$$

and fix $a_i \leq v_{i,t} < r_i$,

$$
\begin{aligned}
(r_i, c_{-i}) &= (R_i(a_{-i}), c_{-i}) \\
&= w_{t+1}(v_{i,t}, a_{-i}),
\end{aligned}
$$

are both incentive compatible continuation values due to implementability of q.

Part (2) I use Corollary 2 to compare $\tilde{q}$ and q.

$$
\begin{aligned}
&\tilde{w}_{i,t} - w_{i,t} \\
&= \mathbb{E}\left[\max\left\{v_{i,t}, \tilde{R}_{i,t}(v_{-i,t})\right\} - \max\left\{v_{i,t}, R_{i,t}(v_{-i,t})\right\} | v_{-i,t} \geq a_{-i}\right] \cdot \mathbb{P}\left\{v_{-i,t} \geq a_{-i}\right\} \\
&\quad + (\tilde{c}_i - c_i) \cdot \mathbb{P}\left\{v_{-i,t} < a_{-i}\right\} \\
&= \mathbb{E}\left[\max\left\{v_{i,t}, r_i\right\} - \max\left\{v_{i,t}, R_{i,t}(v_{-i,t})\right\} | v_{-i,t} \geq a_{-i}\right] \cdot \mathbb{P}\left\{v_{-i,t} \geq a_{-i}\right\} \\
&\geq \mathbb{E}\left[\max\left\{v_{i,t}, r_i\right\} - \max\left\{v_{i,t}, r_i\right\} | v_{-i,t} \geq a_{-i}\right] \cdot \mathbb{P}\left\{v_{-i,t} \geq a_{-i}\right\} \\
&= 0.
\end{aligned}
$$

The inequality is due to the observation Equation 1.5.

Therefore, q is Pareto dominated by $\tilde{q}$. □

Proof of Lemma 4: Consider two other stage mechanisms q^{min} and q^{max} in place of q after history v^{t-1}, where q^{max} is the binary mechanism with thresholds (a_i, r_{-i}) and q^{min} is the binary mechanism with thresholds (a_i, a_{-i}), and with continuation values specified in Figure 1.12. I add superscript *min* and *max* to denote the modified recommendation thresholds,

$$
\begin{aligned}
r_i^{min} &= a_i, \\
r_{-i}^{min} &= r_{-i},
\end{aligned}
$$

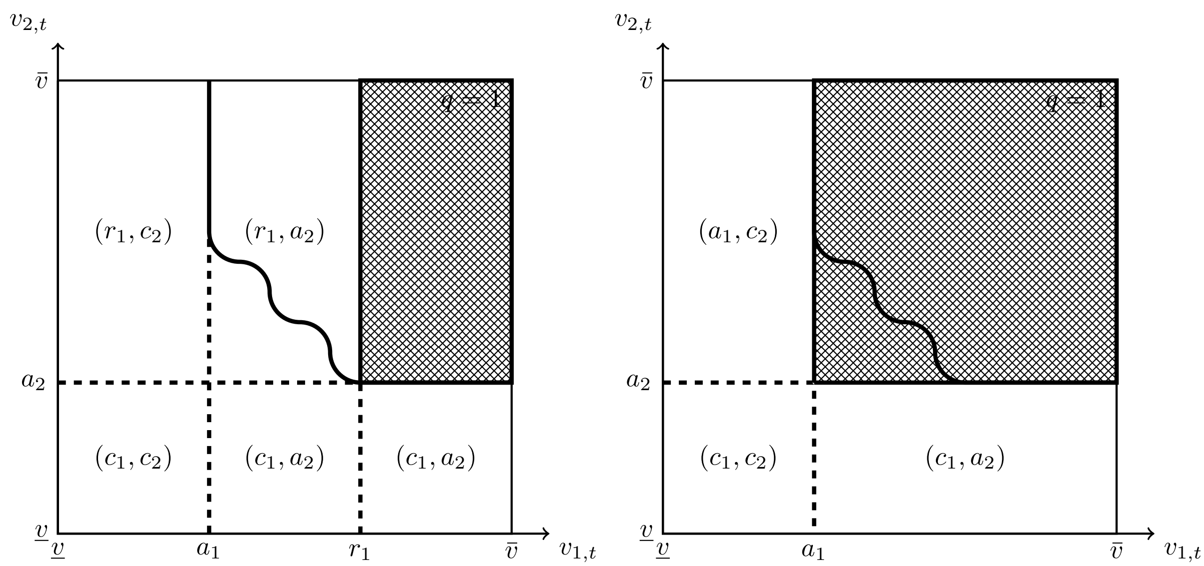

Figure 1.12: Mechanism q^{min}

and,

$$r_i^{max} = a_i,$$

$$r_{-i}^{max} = a_{-i}.$$

with continuation values specified in .

I divide the proof into two parts.

1. q^{min} and q^{max} are incentive compatible.

2. A randomization between Q^{min} and Q^{max} results in the same continuation value for i and is a smaller continuation value for $-i$.

Part (1) The continuation values are incentive compatible for similar reasons as the continuation values of $\tilde{q}$ are incentive compatible in the proof of Lemma 3. The rest follows from Lemma 8.

Part (2) There are three things to show.

1. q^{min} is worse than q for every agent,

2. q^{max} is better than q for i,

3. q^{max} is worse then q for $-i$.

I use Corollary 2 for all three comparisons.

Comparison (1) : For every agent $j \in \{1, 2\}$,

$$
\begin{aligned}
& w_{j,t} - w_{j,t}^{min} \\
& = \mathbb{E}\left[\max\{v_{j,t}, R_{j,t}(v_{-j,t})\} - \max\left\{v_{j,t}, R_{j,t}^{min}(v_{-j,t})\right\} | v_{-j,t} \geq a_{-j}\right] \cdot \mathbb{P}\{v_{-j,t} \geq a_{-j}\} \\
& \quad + \left(c_j^{min} - c_j\right) \cdot \mathbb{P}\{v_{-j,t} < a_{-j}\} \\
& = \mathbb{E}\left[\max\{v_{j,t}, R_{j,t}(v_{-j,t})\} - \max\{v_{j,t}, a_j\} | v_{-j,t} \geq a_{-j}\right] \cdot \mathbb{P}\{v_{-j,t} \geq a_{-j}\} \\
& \geq \mathbb{E}\left[\max\{v_{j,t}, a_j\} - \max\{v_{j,t}, a_j\} v_{-j,t} \geq a_{-j}\right] \cdot \mathbb{P}\{v_{-j,t} \geq a_{-j}\} \\
& = 0.
\end{aligned}
$$

The inequality is due to the observation Equation 1.5.

Comparison (2) : For agent i,

$$
\begin{aligned}
& w_{i,t}^{max} - w_{i,t} \\
& = \mathbb{E}\left[\max\left\{v_{i,t}, R_{i,t}^{max}(v_{-i,t})\right\} - \max\{v_{i,t}, R_{i,t}(v_{-i,t})\} | v_{-i,t} \geq a_{-i}\right] \cdot \mathbb{P}\{v_{-i,t} \geq a_{-i}\} \\
& \quad + (c_i^{max} - c_i) \cdot \mathbb{P}\{v_{-i,t} < a_{-i}\} \\
& = \mathbb{E}\left[\max\{v_{i,t}, r_i\} - \max\{v_{i,t}, R_{i,t}(v_{-i,t})\} | v_{-i,t} \geq a_{-i}\right] \cdot \mathbb{P}\{v_{-i,t} \geq a_{-i}\} \\
& \geq \mathbb{E}\left[\max\{v_{i,t}, r_i\} - \max\{v_{i,t}, r_i\} | v_{-i,t} \geq a_{-i}\right] \cdot \mathbb{P}\{v_{-i,t} \geq a_{-i}\} \\
& = 0.
\end{aligned}
$$

The inequality is due to the observation Equation 1.5.

Comparison (3) : For agent $-i$,

$$
\begin{aligned}
& w_{-i,t} - w^{max_{-i,t}} \\
& = \mathbb{E}\left[\max\{v_{-i,t}, R_{-i,t}(v_{i,t})\} - \max\{v_{-i,t}, R^{max_{-i,t}}(v_{i,t})\} | v_{i,t} \geq a_i\right] \cdot \mathbb{P}\{v_{i,t} \geq a_i\} \\
& \quad + \left(c_j^{max} - c_j\right) \cdot \mathbb{P}\{v_{i,t} < a_i\} \\
& = \mathbb{E}\left[\max\{v_{j,t}, R_{j,t}(v_{-j,t})\} - a_j | a_i \leq v_{i,t} < r_i\right] \cdot \mathbb{P}\{a_i \leq v_{i,t} < r_i\} \\
& \geq \mathbb{E}\left[\max\{v_{j,t}, a_j\} - a_j | a_i \leq v_{i,t} < r_i\right] \cdot \mathbb{P}\{a_i \leq v_{i,t} < r_i\} \\
& \geq 0.
\end{aligned}
$$

The first inequality is due to the observation Equation 1.5. Note that this abuses Corollary 2 to since the non-threshold regions are not the same.

Therefore, there is a randomization between q^{min} and q^{max} such that the continuation value for i is the same and lower for agent $-i$. $\square$

Proof of Theorem 1: Since binary mechanisms in Lemma 4 are ternary and the set $-i$ is a singleton, the result follows from Lemma 4 and Lemma 3. $\square$

Chapter 2

School Allocation Problem with Observable Characteristics

2.1 Introduction

The literature on school choice studies the problem of how to allocate students into schools. A typical problem includes students with different preferences, and possibly schools with different preferences or priorities for students.

One approach to the school choice problem was introduced in Bogomolnaia and Moulin [8]. They consider a mechanism design approach. Students report their preferences and the mechanism allocates the students into the schools. A good mechanism satisfies certain properties such as the following,

1. Efficiency: the resulting allocation is not Pareto dominated by any other feasible allocation.

2. Envy-freeness: the resulting allocation for any individual student is not dominated by the allocation for another student.

3. Symmetry: the resulting allocation is the same for students with the same preferences.

Bogomolnaia and Moulin [8] introduce a mechanism called the simultaneous eating mechanism that is characterized by property (1). They also show that a special case of their simultaneous eating mechanism, called the probabilistic serial mechanism, satisfies all three properties $(1),(2)$ and (3). Later, Liu and Pycia [26] shows that for problems with full support - meaning that for every possible permutation of the preference ordering over the schools, a strictly positive proportion of students have that preference - a mechanism satisfies $(1),(2)$ and (3) if and only if it is probabilistic serial.

The symmetry assumption made by Bogomolnaia and Moulin [8] has an important interpretation as fairness criterion: no two students are treated differently. It also has a practical dimension: if the only information about students comes from their reports, only the reports can be taken into account by the mechanism. However, the assumption is not appropriate for many contexts. In many situations, categories of students can be observed by the schooling authority. Examples include,

- The further the location of a student's home to a school, the higher the transportation cost and as a result, ceteris paribus, the less the student will prefer to go that school. The school board may have additional data on the average amount of saving on transportation costs for each student, and as a result have good estimate of the cardinal welfare for every feasible allocation.

- Students with higher standardized test grades may be more likely to enjoy a school with larger libraries, more labs and better teachers. The students may not be able to evaluate the benefit from better libraries, but the school board is able to provide a more accurate estimate of the difference in benefit to students with different test scores.

- The more siblings a student has in a school, the higher the utility the student will obtain when attending that school. In this case, two students with the same ordinal preference ordering may have dramatically different cardinal utilities from attending the schools. It is impossible to obtain from the students the cardinal utility gain from having a sibling in the same school because there are no consistent units of measurement; however, the school board may be able to evaluate the gain relatively consistently.

- Students with disabilities may benefit more from schools with more specialized facilities, equipment and teachers with training for students with special needs. Similarly to the previous examples, the school board has better information on the utility gain for these students.

In the above situations, there are important reasons to treat different students differently. For instance, the school authority may want to reduce the transportation costs and give the priority to the students who live in the school's neighborhood. Similarly, the priority to schools that are accessible can be given to students with disabilities.

In this paper, I consider a version of the school choice problem where students that belong to different well-defined groups can be treated differently. Formally, as in Bogomolnaia and Moulin [8], I take the mechanism design approach, and I consider stochastic allocations and compare them through first-order stochastic dominance. I will make the assumption that although the students have cardinal utilities, they can only report ordinal rankings, and as a result, the mechanism is ordinal. For tractability, I also assume continuum of students of mass 1.

I keep the efficiency assumption and replace envy-freeness and symmetry with a pair of weaker assumptions that students with the same characteristic do not envy each other's allocations and they must be treated equally and call these assumptions within-group envy-freeness and within-group symmetry. Consequently, I consider the class of mechanisms that satisfy the following properties:

1. Efficiency: the resulting allocation is not Pareto dominated by any other feasible allocation

2. Within-group envy-freeness: the resulting allocation for any individual student is not dominated by the allocation for another student *with the same characteristics* .

3. Within-group symmetry: the resulting allocation is the same for students *with the same characteristics* with the same preferences.

The main result is similar to the one in Liu and Pycia [26] that the only mechanisms that satisfy the above three properties are from a class of modified versions of the probabilistic serial mechanisms from Bogomolnaia and Moulin [8]. These mechanisms are probabilistic serial with school-specific subcapacities for each group of students with the same characteristic. Each subcapacity specifies the maximum amount of students with a particular characteristic that are allowed to be assigned to the school. In addition, given data on the expected cardinal utilities from assigning students with a particular characteristic to a specific school, I show that the cardinally efficient allocations can also be obtained from the modified probabilistic serial mechanisms with subcapacities that can be chosen according the solution to a convex programming problem.

Here, I briefly describe the algorithm as if the students are consuming portions of schools continuously: for students in a specific group, schools start with sizes equal to their subcapacities and every student starts eating her favorite school at the same unit rate until some schools are completely eaten. Then, students start to eat their favorite among the remaining schools that are not eaten. The process ends at time 1, and the total amount of school a student has eaten will represent the probability that she is allocated to school.

This paper is most closely related to the literature on ordinal mechanisms for large markets developed by Kojima and Manea [22], and Che and Kojima [9] and extended by Liu and Pycia [26]. The proofs also use techniques from these papers. The literature attempts to find efficient, strategy-proof and symmetric mechanisms for large markets; however, the focus of this paper deviates from the literature in that it studies problems in which the students have observable characteristics that can provide additional information for the designers. As a result, this paper partially addresses the loss of welfare due to the restriction to ordinal mechanisms mentioned in Abdulkadiroğlu, Che, and Yasuda [1] and Pycia [35] by proposing the use of estimated cardinal utilities (from students' observable characteristics) in addition to the elicited ordinal preference. The results are not directly applicable to finite markets, but the algorithm in this paper can be applied to finite markets and potentially improve welfare in high school choice assignments in Boston or New York described in Pathak and Sethuraman [31], Abdulkadiroğlu, Pathak, and Roth [2] and Pathak and Sönmez [32].

Section 2.2 sets up the model with ordinal preferences. Section 2.3 defines the class of modified probabilistic serial mechanisms with subcapacities and explains why they are the only mechanisms that are efficient, within-group envy-free and symmetric. Section 2.4 sets up the model for the planner with cardinal preferences and explains the convex programming problem to solve for the optimal subcapacities.

2.2 Ordinal Model

In this section, I describe the model.

There is a finite set of schools $L = \left\{1, 2, ..., \bar{L}\right\}$ and a finite set of observable characteristics $K = \left\{1, 2, ..., \bar{K}\right\}$. Each student has one of these characteristics, thus the students are partitioned into groups by their characteristics. School l has mass of c_l spots and there are mass μ_k of students with characteristic k. The total mass of the students is 1 and I assume that there is enough room in the schools to assign each student,

$$\sum_{l \in L} c_l \geq \sum_{k \in K} \mu_k = 1.$$

A student's preference is given by a strict ordering $p \in \mathcal{P}(L)$, where $\mathcal{P}(L)$ is the set of all permutations (strict orderings) of L. A *type* is a pair (k, p), and $\mu(k, p)$ denotes the mass of students in group k who have preference ordering p. Let the individual students be indexed by $i \in \mathcal{I} = [0, 1]$, equipped with the Lebesgue measure λ.

Definition 7. A *(full support) profile* of students is a λ-measurable function, *type* $: \mathcal{I} \to K \times \mathcal{P}(L)$, that satisfies the following conditions:

1. Type and group consistency,

$$\sum_{p \in \mathcal{P}(L)} \mu(k, p) = \mu_k,$$

where,

$$\mu(k,p) = \int_{i\in\mathcal{I},type(i)=(k,p)} d\lambda.$$

2. Full support,

$$\mu(k,p) > 0 \, \forall \, (k,p) \in K \times \mathcal{P}(L).$$

In this definition, a profile maps the students' names to their types, (1) states that the total mass of students with type (k,p) is $\mu(k,p)$ and the total mass of students in group k is μ_k, and (2) states that the measure of students with every possible type is strictly positive. More importantly, each profile induces a full support distribution of types, μ, where $\mu(k,p)$ is the mass of students of type (k,p).

Definition 8. An *allocation* is a measurable function $q : \mathcal{I} \to \Delta L$, where $q\,(l;i)$ is the probability that a student $i \in \mathcal{I}$ is assigned to school l. An allocation q is *feasible* if,

$$\int_{i\in\mathcal{I}} q\,(l;i)\, d\lambda \leq c_l.$$

An allocation maps each student in $\mathcal{I}$ to a stochastic allocation. One interpretation is that students receive their spot randomly according the distribution $q\,(i)$. Each stochastic allocation can be implemented by a randomized mechanism. For fixed μ_k, let $\mathcal{M}$ be the set of all profiles and $\mathcal{Q}$ be the set of all feasible allocations.

Definition 9. A *mechanism* is a function $Q : \mathcal{M} \to \mathcal{Q}$, that maps every profile to a feasible allocation.

Next, I define ordinal efficiency using first order stochastic dominance.

Definition 10. An allocation q is *dominated* by q' for student i with preference p if

$$\sum_{l'\succ_p l} q\,(l';i) \leq \sum_{l'\succ_p l} q'\,(l';i) \;\; \forall \, l \in L,$$

with strict inequality for at least one l.

One allocation dominates another if it first order stochastically dominates the other allocation. The three desired properties of an ordinal mechanism in this model are the following. I will use the notation $k : \mathcal{I} \to K$ to represent the function that maps a student to her group characteristics.

Definition 11. Efficiency: $q\,(i)$ is not dominated by any $q'\,(i)$ for any $i \in \mathcal{I}$.

Definition 12. Within-group envy-free: $q(i)$ is not dominated by any $q(j)$ for any $i, j \in \mathcal{I}$ such that $k(i) = k(j)$.

Definition 13. Within-group symmetry: for each type (k,p), $q(i) = q(j)$ for any $i, j \in \mathcal{I}$ such that $type(i) = type(j) = (k,p)$.

Within-group symmetry states that students with the same characteristics and preferences should be assigned the same allocation. Given this assumption, I can use the notation $q(l;k,p)$ to denote the probability of a student with type (k,p) getting allocated the school l. Envy-freeness states that any a student (k,p) will not prefer the allocation of another student (k,p') for $p' \neq p$. The assumption is a weaker version of strategy-proofness. Efficiency implies that no other allocation is preferred by all students.

In the case in which every student with the same type gets the same allocation, the notation $q(l;k,p)$ will be used in place of $q(l;i)$ to denote the probability that a student with type (k,p) is assigned to school l. Consequently, the total amount of students with type (k,p) who are assigned to school l satisfies,

$$\mu(k,p) \cdot q(l;k,p) = \int_{i \in \mathcal{I}: type(i)=(k,p)} q(l;i)\, d\lambda.$$

2.3 Modified Probabilistic Serial

In this section, I describe the probabilistic serial mechanism from Bogomolnaia and Moulin [8] and the modification to include subcapacities that represent the maximum number of students with a certain characteristic that can be allocated to each school. I also state and explain the main result that a mechanism is efficient, within-group envy-free and within-group symmetric if and only if it is modified probabilistic serial.

I start by briefly summarizing the mechanism in Bogomolnaia and Moulin [8]. I describe their algorithm as if the students are consuming portions of schools continuously at a fixed rate. Schools start with sizes equal to their capacities and every student starts eating her favorite school at rate 1 until some schools are completely eaten. Then, students start to eat their favorite among the remaining schools that are not eaten. The process ends at time 1, and the total amount of school a student has eaten will represent the probability that she is allocated to school.

Then I define the collection of subcapacities $\{c_l^k\}_{k \in K, l \in L}$. Subcapacity c_l^k represents the maximum number of students in group k that can be assigned to school l by the algorithm. They must satisfy the feasibility condition,

$$\sum_{k \in K} c_l^k \leq c_l \; \forall \, l \in L.$$

I modify the mechanism in Bogomolnaia and Moulin [8] in two ways.

1. Schools start with sizes equal to their capacities but are divided into subcapacities for each group specified by $\left\{c_l^k\right\}_{l \in L, k \in K}$, so the initial sizes of the schools are not equal to their actual capacities.

2. Students in one group cannot eat the portion of the schools allocated to other groups.

In this algorithm, all students with the same type get the same allocation. Therefore, within-group symmetry is ensured. The formal descriptions of the algorithms are as follows. Before that, I define the following function.

$M(l, A)$ is a function that maps a school, l, and a set of available schools, $A \subseteq L$, to the set of preference orderings in which the favorite school among the set A is l,

$$M(l, A) = \left\{p \in \mathcal{P}(L) : l \succ l' \; \forall \, l' \in A \setminus \{l\}\right\}. \tag{2.1}$$

Algorithm 1. *Given subcapacities* $\left\{c_l^k\right\}_{l \in L, k \in K}$*, the probabilistic serial mechanism (PS) assigns to each profile the allocation resulting from the following process, assuming* μ *is the distribution of types induced by the given profile.*

Initialize: $L_k^0 = L, y_k^0 = 0$ *for each* $k \in K$ *and* $q^0(l; k, p) = 0$ *for each* $l \in L, k \in K, p \in \mathcal{P}(L)$*,*

Iteration: Assume $L_k^{s-1}, y_k^{s-1}, q^{s-1}$ *are defined for each* k*. For each* k*,*

1. *For each* l*, find* $y_k^s(l)$*, the earliest time at which students in group* k *finish consuming school* l*,*

$$y_k^s(l) = \arg\min_y \left\{ \sum_{p \in M\left(l, L_k^{s-1}\right)} \mu(k, p) \left(y - y_k^{s-1}\right) + \sum_{p \in \mathcal{P}(L)} \mu(k, p)\, q^{s-1}(l; k, p) = c_l^k \right\}.$$

2. *Find* y_k^s*, the earliest time at which students in group* k *finish consuming any school,*

$$y_k^s = \min_l y_k^s(l).$$

3. *Find* F_k^s*, the set of schools that are completely consumed by students in group* k*,*

$$F_k^s = \arg\min_l y_k^s(l).$$

4. *Find the remaining set of available schools for students in group k,*

$$L_k^s = L_k^{s-1} \setminus F_{k.}^s$$

5. *Find $q^s(l;k,p)$, the temporary allocation of school l for students with type (k,p), which represents the amount of school l the students have eaten so far until the end of step s,*

$$q^s(l;k,p) = q^{s-1}(l;k,p) + \mathbb{1}_{p \in M\left(l, L_k^{s-1}\right)}\left(y^s - y^{s-1}\right).$$

For a full support profile, Bogomolnaia and Moulin [8] showed that probabilistic serial generates an allocation is efficient, envy-free, symmetric, and Liu and Pycia [26] showed that an allocation is efficient, envy-free, and symmetric if and only if it is generated by probabilistic serial. The result can be extended to the problem with multiple groups with a similar proof to Theorem 1 in Liu and Pycia [26].

Proposition 1. *An allocation, q, is efficient, within-group symmetric, and envy-free for a full support profile that induces type distribution μ if and only if it is generated by Algorithm 1 modified probabilistic serial with subcapacities,*

$$c_l^k(\mu) = \sum_{p \in \mathcal{P}(L)} \mu(k,p)\, q(l;k,p).$$

The above formula for c_l^k says nothing about how to choose the subcapacities in practice since they are just one of many capacities that are compatible with the allocation q.

In particular, note that ordinal efficiency does not rely on fixing the subcapacities. The probabilistic serial allocation is unconstrained ordinally efficient among all allocations associated with all possible choices of subcapacities.

2.4 Cardinal Model

In this section, I introduce the model with cardinal utility functions that is ordinally consistent with the previous model. I explain that the cardinally efficient allocation can be obtained from the modified probabilistic serial mechanism with subcapacities which can be found as the solution to a convex (linear) optimization problem.

The cardinal utility functions should be compatible with the preference relations from ordinal model. Let $u(i) : L \to \mathbb{R}_+$ be the utility function of a student i, where $u(l;i)$ represents the utility from attending school l. A utility function induces a preference relation p if for every student i with the preference p,

$$l \succ_p l' \text{ whenever } u\left(l;i\right) > u\left(l';i\right) \ \forall\, l \in L.$$

Here, I do not need to assume the students with the same type (k,p) have the same utility function, but since the within-group symmetry assumption requires the allocation to be the same for students with the same type, and the welfare function uses the average utilities, I can restrict attention to using the average utilities over the same type,

$$u\left(l;k,p\right) = \frac{1}{\mu(k,p)} \int_{type(i)=(k,p),p(u(i))=p} u\left(l;i\right) d\lambda,$$

where $p\left(u\right)$ is the preference relation induced by the utility ranking u.

Definition 14. A utility distribution that is *consistent* with a preference profile μ is the set of utility functions $u\left(i\right)$, such that,

$$\int_{type(i)=(k,p),p(u(i))=p} d\lambda = \mu(k,p)\ \forall\ k \in K, p \in \mathcal{P}(L).$$

Definition 15. The allocation $q^\star$ is *cardinally efficient* if it maximizes the average expected welfare:

$$q^\star \in \arg\max_{q \in \mathcal{Q}} \int_{i \in \mathcal{I}} \sum_{l \in L} u\left(l;i\right) \cdot q\left(l;i\right) d\lambda.$$

If the mechanism is restricted to be within-group symmetric, then $q^\star$ maximizes the welfare function

$$W\left(q\right) = \sum_{(k,p)\in K\times\mathcal{P}(L)} \sum_{l\in L} u\left(l;k,p\right) \cdot q\left(l;k,p\right) \cdot \mu(k,p).$$

Since probabilistic serial is ordinally efficient, any cardinally efficient allocation must be obtained by probabilistic serial for some subcapacities. Therefore, I write the allocation and welfare as a function of the subcapacities in the probabilistic serial mechanism that generates them, let $c = \left\{c_l^k\right\}_{l\in L,k\in K}, c^k = \left\{c_l^k\right\}_{l\in L}$,

$$W_k\left(c^k\right) = \sum_{l\in L} u\left(l;k,p\right) q^{\text{ PS }(c)}\left(l;k,p\right)\mu(k,p),$$

and,

$$W(c) = \sum_{k \in K} W_k\left(c^k\right),$$

where $q^{\mathrm{PS}\,(c)}$ is the allocation generated by probabilistic serial with subcapacities c.

Proposition 2. *The function $W(c)$ is concave in c.*

In order to find welfare-maximizing mechanism, I need to find the maximum of function $W(c)$ subject to the following linear constraints.

$$\begin{aligned}
&\max_{c_l^k \in [0, c_l]} W(c) \\
&\text{such that } \sum_{k \in K} c_l^k = c_l \\
&\text{and } \sum_{l \in L} c_l^k = \mu_k
\end{aligned}$$

The first constraint is the school capacity constraint requiring the subcapacities for a school for all groups add up to the total physical capacity of the school. The second constraint is the profile constraint requiring the subcapacities for a group from all schools add up to the total mass of the students in that group. Both constraints hold at optimum because of the assumption that there is enough room in the schools to allocate all the students.

The optimization can be done computationally. Proposition 2 says that function $W(c)$ is concave in c, hence, standard gradient methods can be used to find the maximum. One conjecture is that the choice of c can be found easily using the differed acceptance algorithm. It is indeed the case when the number of schools is two, and it can be shown by enumerating all possible cases. I am unable to verify at the moment that the conjecture is correct for problems with more than two schools as there are no simple characterization of the allocations resulted from probabilistic serial. I leave this interesting question for future research.

2.5 Proofs

2.5.1 Proof of Proposition 1

I divide the proof into three parts, stated as the following three lemmas.

Lemma 9. *Modified probabilistic serial is (ordinally) efficient.*

Lemma 10. *Modified probabilistic serial is envy-free.*

Lemma 11. *Given a full-support profile, μ, if an allocation q is efficient, within-group envy-free and symmetric, then it is generated by modified probabilistic serial with constraints $c_l^k = \sum_{p \in \mathcal{P}(L)} q(l;k,p)\,\mu(k,p)$.*

Lemma 9 and Lemma 10 are modified from Theorem 1 and Proposition 1 of Bogomolnaia and Moulin [8], respectively, and Lemma 11 is modified from Theorem 1 of Liu and Pycia [26]. The modified probabilistic serial mechanism is within-group symmetric by construction, so these three lemmas, together with the full support assumption, implies Proposition 1.

Proof of Lemma 9: Suppose, for a contradiction that q is obtained by modified probabilistic serial and it is not efficient, and q is dominated by q'.

Let (k, p_1) be the student such that $q(k,p_1) \neq q'(k,p_1)$, there are l_0, l_1 with $l_0 \succ_{p_1} l_1$, such that,

$$q(l_1;k,p_1) > q'(l_1;k,p_1),$$
$$q(l_0;k,p_1) < q'(l_0;k,p_1).$$

Then $l_0 \succ_{p_1} l_1$ and $q(l_1;k,p_1) > 0$.

Similarly, there is $l_1 \succ_{p_2} l_2$ and $q(l_2;k,p_2) > 0$.

Since L is finite, there exists some cycle,

$$l_0 \succ_{p_1} l_1 ... l_R \succ_{p_R} l_{0,}$$

such that for every $r \in \{0, 1, ..., R\}$,

$$l_{r-1} \succ_{p_r} l_r \text{ and } q(l_r;k,p_r) > 0.$$

Take an arbitrary r, let the q^s represent the partial probabilistic serial allocation at time s, and define the following,

$$s_r = \inf_s \{s : q^s(l_r;k_r,p_r) > 0\}.$$

Note that $p_{r-1} \notin L_{k_r}^{s_{r-1}}$, where the notation L_k^s is introduced in step (4) of Algorithm 1 as the remaining set of available schools for students in group k after step s. This implies $s_{r-1} < s_{r.}$

This is true for for every r, implying $s_0 < s_1 < ... < s_{R-1} < s_0$, which leads to a contradiction. □

Proof of Lemma 10: Fix a student (k,p) with $l_1 \succ_p l_2 ... \succ_p l_{\bar{L}}$, and let s_1 be the earliest time at which l_1 is completely eaten, meaning,

$$l_1 \in L_k^{s_1-1} \setminus L_k^{s_1}.$$

For $s \leq s_1 - 1, (k,p) \in M\left(l_1, L_k^s\right)$, where $M\left(l, A\right)$ is introduced before Algorithm 1 by Equation 2.1 as the set of preference orderings in which the favorite school among the set A is l, and

$$q^{s_1}\left(l_1; k, p\right) = y_k^{s_1} \geq q^{s_1}\left(l_1; k, p'\right) \ \forall \ p' \neq p.$$

Since l_1 is completely eaten at s_1, the previous relationship holds for final allocation q as well,

$$q\left(l_1; k, p\right) \geq q\left(l_1; k, p'\right) \ \forall \ p' \neq p.$$

Now, let s_2 be the earliest time $\{l_1, l_2\}$ is completely eaten, $s_2 \geq s_1$ and for the same reason,

$$q\left(l_1; k, p\right) + q\left(l_2; k, p\right) = y_k^{s_2} \geq q\left(l_1; k, p'\right) + q\left(l_2; k, p'\right) \ \forall \ p' \neq p.$$

Repeat the argument to see that $q\left(k,p\right)$ dominates $q\left(k,p'\right)$ for every $p' \neq p$. □

Proof of Lemma 11: Fix any allocation, q', that is efficient, within-group envy-free and symmetric, and the allocation, q^1, obtained by modified probabilistic serial. Let q^t denote the partial allocation at time $t \in [0,1]$ from the probabilistic serial.

I show that for any student with type $(k,p) \in K \times \mathcal{P}(L)$, any school $l \in L$, and at any time $t \in [0,1]$, q' dominates q^t. Then, by the efficiency property of q^1 from Lemma 9, $q' = q^1$, which concludes the proof.

Assume for a contradiction, let τ be the earliest time when q' does not dominate q^t. Using the notations in Definition 10, for some school l and and student with type (k,p),

$$\tau = \inf \left\{ t : \sum_{l' \succ_p l} q'\left(l'; k, p\right) < \sum_{l' \succ_p l} q^t\left(l'; k, p\right) \right\}.$$

By continuity of the function q^t in t, q' dominates q^t for each $t \in [0, \tau]$. In particular, at time τ, the students with type (k,p) must be eating some school l, and,

$$\sum_{l' \succ_p l} q'\left(l'; k, p\right) \geq \sum_{l' \succ_p l} q^\tau\left(l'; k, p\right) = \tau.$$

Since τ the earliest time the above inequality stops holding,

$$\sum_{l' \succ_{k,p} l} q'(l'; k, p) = \sum_{l' \succ_{k,p} l} q^{\tau}(l'; k, p) = \tau.$$

Now, using the full support condition, l must be the favorite object of some agent (k, p'), and,

$$q'(l; k, p') \geq q^{\tau}(l; k, p') = \tau.$$

The envy-free assumption implies (k, p) does not prefer the allocation of (k, p'),

$$q'(l; k, p') \leq \tau.$$

Therefore,

$$q'(l; k, p') = \tau.$$

Since school l is not completely eaten at time τ,

$$q^1(l; k, p') > \tau.$$

Therefore, the equalities imply that (k, p') gets less l in q' than q^1. The efficiency of q^1 implies that there is another student $(k, \hat{p})$ who gets,

$$q'(l; k, \hat{p}) > q^1(l; k, \hat{p}).$$

And there is some school $\hat{l}$ that is not l that student $(k, \hat{p})$ prefers just more than school l,

$$\sum_{l' \succ_{\hat{p}} \hat{l}} q'(l'; k, \hat{p}) \geq \sum_{l' \succ_{\hat{p}} \hat{l}} q^{\tau}(l'; k, \hat{p}),$$

implying,

$$\sum_{l' \succ_{\hat{p}} \hat{l}} q'(l'; k, \hat{p}) \geq \tau - q^{\tau}(l; k, \hat{p}),$$

and,

$$\sum_{l' \succ_{\hat{p}} \hat{l}} q'(l'; k, \hat{p}) \geq \tau - q^1(l; k, \hat{p}),$$

and finally,

$$\sum_{l' \succ_{\hat{p}} l} q'(l'; k, \hat{p}) > \tau.$$

Comparing students (k, p') and $(k, \hat{p})$, envy-freeness of q' leads to a contradiction. □

2.5.2 Proof of Proposition 2

I prove several smaller lemmas about properties of envy-free allocations in order to prove concavity.

Lemma 12. *Any convex combination of envy-free allocations is envy-free.*

Lemma 13. *Any inefficient envy-free allocation has an envy-free Pareto improvement.*

Lemma 14. *The set of envy-free allocations are closed.*

Proof of Lemma 12: Consider arbitrary pair of students (k, p) and (k, p') under two different allocations q_1 and q_2.

Let $p = l_1 \succ l_2 \succ l_3 ... \succ l_{\bar{L}}$ be the preference ranking of the first student, and define the following,

$$t_1^i = \min_t \left\{ \sum_{s=1}^{t} q_1(l_s; k, p) < \sum_{s=1}^{t} q_1(l_s; k, p') \right\}$$
$$t_1^a = \max_t \left\{ \sum_{s=1}^{t} q_1(l_s; k, p) \leq \sum_{s=1}^{t} q_1(l_s; k, p') \right\}$$
$$t_2^i = \min_t \left\{ \sum_{s=1}^{t} q_2(l_s; k, p) < \sum_{s=1}^{t} q_2(l_s; k, p') \right\}$$
$$t_2^a = \max_t \left\{ \sum_{s=1}^{t} q_2(l_s; k, p) \leq \sum_{s=1}^{t} q_2(l_s; k, p') \right\}$$

By envy-freeness, $t_1^i \neq t_1^a$ and $t_2^i \neq t_{2,}^a$

Consider a convex combination $q_0 = (\alpha) q_1 + (1 - \alpha) q_2$ for $\alpha \in [0, 1]$, For $t \leq \min\left\{t_1^i, t_2^i\right\}$,

$$\begin{aligned}\sum_{s=1}^{t} q_0\left(l_s;k,p\right) &= \sum_{s=1}^{t}\left(\alpha\right) q_1\left(l_s;k,p\right) + \left(1-\alpha\right) q_2\left(l_s;k,p\right) \\ &< \sum_{s=1}^{t}\left(\alpha\right) q_1\left(l_s;k,p'\right) + \left(1-\alpha\right) q_2\left(l_s;k,p'\right) \\ &= \sum_{s=1}^{t} q_0\left(l_s;k,p'\right)\end{aligned}$$

And for $t \geq \max\left\{t_1^a, t_2^a\right\}$,

$$\begin{aligned}\sum_{s=1}^{t} q_0\left(l_s;k,p\right) &= \sum_{s=1}^{t}\left(\alpha\right) q_1\left(l_s;k,p\right) + \left(1-\alpha\right) q_2\left(l_s;k,p\right) \\ &> \sum_{s=1}^{t}\left(\alpha\right) q_1\left(l_s;k,p'\right) + \left(1-\alpha\right) q_2\left(l_s;k,p'\right) \\ &= \sum_{s=1}^{t} q_0\left(l_s;k,p'\right)\end{aligned}$$

Therefore, under q_0, no student strictly prefers the allocation of another student, q_0 is envy-free. □

Proof of Lemma 13: Consider an allocation q and another allocation q' that (Pareto) dominates q.

For each student (k,p) and pair of schools l_1 and l_2, define the flow from school l_1 to l_2 by $\Delta\left(l_1,l_2;k,p\right)$ satisfying:

$$\begin{aligned}\sum_{l_2\in L}\Delta\left(l_1,l_2;k,p\right) &= \max\left\{0, q\left(l_1;k,p\right) - q'\left(l_1;k,p\right)\right\}, \\ \sum_{l_1\in L}\Delta\left(l_1,l_2;k,p\right) &= \max\left\{0, q\left(l_2;k,p\right) - q'\left(l_2;k,p\right)\right\}, \\ \Delta\left(l_1,l_2;k,p\right) &\geq 0.\end{aligned}$$

Then define another allocation $q^\star$ by:

$$\begin{aligned}q^\star\left(l_1;k,p\right) = q\left(l_1;k,p\right) &- \sum_{l_2\in L}\mathbb{1}_{\Delta\left(l_1,l_2;k,p\right)>0 \text{ or } l_2\succ_p l_1}\cdot\Delta^\star\left(l_1,l_2;k,p\right) \\ &+ \sum_{l_2\in L}\mathbb{1}_{\Delta\left(l_2,l_1;k,p\right)>0 \text{ or } l_2\succ_p l_1}\cdot\Delta^\star\left(l_2,l_1;k,p\right)\end{aligned}$$

where $\Delta^\star$ is defined as:

$$\Delta^\star(l_1,l_2;k,p)=\frac{\sum_{(k',p')\in K\times\mathcal{P}(L)}\Delta(l_2,l_1;k',p')\cdot\mu(k',p')}{\sum_{(k',p'):l_1\succ_{p'}l_2 \text{ and } \Delta(l_1,l_2;k',p')=0}\mu(k',p')+\sum_{(k',p')\in K\times\mathcal{P}(L)}\Delta(l_2,l_1;k',p')\cdot\mu(k',p')}$$

Note that the flows from q to q' and the flows from q to $q^\star$ are the same since the previous system for Δ is still satisfied.

Also, $q^\star$ still dominates q since,

$$\begin{cases}\Delta^\star(l_1,l_2;k,p)>0 & \text{if } l_1\succ_p l_2\\ \Delta^\star(l_1,l_2;k,p)<\Delta(l_1,l_2;k,p) & \text{if } l_2\succ_p l_1\end{cases}$$

And $q^\star$ is envy-free since,

$$\begin{cases}\Delta^\star(l_1,l_2;k,p)\geq\Delta^\star(l_1,l_2;k,p') \text{ for every } p' & \text{if } l_1\succ_p l_2\\ \Delta^\star(l_1,l_2;k,p)\geq 0 & \text{if } l_2\succ_p l_1\end{cases}$$

Therefore, $q^\star$ is an envy-free Pareto improvement to q. □

Proof of Lemma 14: Consider any sequence of allocations $\{q_n\}_{n=1}^{\infty}$ and the element-wise limit $q^\star$.

Fix any two students (k,p) and (k,p'), since q_n are envy-free for each n,

$$\sum_{l\succ_p l'}q_n(l;k,p')\leq\sum_{l\succ_p l'}q_n(l;k,p)\ \forall\, l'\in L,$$

where l_s is the s -th school in the preference ranking of student (k,p).

Then,

$$\lim_{n\to\infty}\sum_{l\succ_p l'}q_n(l;k,p')\leq\lim_{n\to\infty}\sum_{l\succ_p l'}q_n(l;k,p)\ \forall\, l'\in L,$$
$$\sum_{l\succ_p l'}q^\star(l;k,p')\leq\sum_{l\succ_p l'}q^\star(l;k,p)\ \forall\, l'\in L.$$

Therefore, $q^\star$ is envy-free. The set is closed under limits.

Similarly, the set of Pareto improvements of any allocation is closed. □

Proof of Proposition 2: Let c,c' be two vectors of subcapacities, and q,q' be the probabilistic allocation with subcapacities c,c' respectively.

Consider allocation $q_0 = (\alpha)\, q + (1-\alpha)\, q'$ and the welfare of allocation q_0 is $(\alpha)\, W(c) + (1-\alpha)\, W(c')$

If q_0 can be obtained from PS with capacities $(\alpha)\, c + (1-\alpha)\, c'$, then $(\alpha)\, W(c) + (1-\alpha)\, W(c') = W((\alpha)\, c + (1-\alpha)\, c')$.

Suppose q_0 is obtained from probabilistic serial, and since q_0 is envy-free from Lemma 12, q_0 is not efficient by Proposition 1.

Let V be the set of envy-free allocations that Pareto dominates $q_{0.}$

V is bounded since the set of allocations is bounded and the set of all envy-free allocations and the set of allocations that are Pareto improvements to q_0 are closed by Lemma 8. Then, V is an intersection of two compact sets implying that V is compact.

Therefore, there is an allocation $q^\star \in V$ that maximizes $W(\cdot)$.

Note that $q^\star$ must be efficient because if not, by Lemma 13, there is a envy-free Pareto improvement of $q^\star$ in V which contradicts the definition that $q^\star$ maximizes $W(\cdot)$.

$q^\star$ is envy-free and efficient, implying that $q^\star$ is the probabilistic serial allocation with capacity $(\alpha)\, c + (1-\alpha)\, c'$.

Therefore, $(\alpha)\, W(c) + (1-\alpha)\, W(c') \leq W((\alpha)\, c + (1-\alpha)\, c')\,, W$ is concave in c.

The function W(c) is non-decreasing for c_l^k due to the assumption that $u(\cdot) \geq 0$. $\square$

Chapter 3

Mechanism Design for Stopping Problems with Two Actions

3.1 Introduction

In many economic situations, an agent wants to choose between two alternatives but is not required to make the choice right away. Often, the agent can postpone the choice until later in order to either gather more information or wait for more favorable conditions. Additionally, another party, called a principal, has an interest in the timing and choices made by the agent and can use transfers to provide incentives.

I model such a situation as a stopping problem with two actions. The agent observes a Markov state, decides when to stop, and upon stopping, chooses between one of the two actions. The agent's payoffs are a sum of an intrinsic payoff, which depend on the time, action, state, and a quasi-linear transfer received from the principal, which depends on the time and the action choice. The goal of this analysis is to describe the range of choice rules that can be implemented and to determine whether and how the principal can incentivize the agent with quasi-linear transfers.

Economic applications include:

- (Job Search) A worker observes a changing demand for his services and chooses between taking up a job, continuing to search, or leaving the market (possibly to retire, or to go back to school). The government decides on the amount of employment insurance it offers to workers in order to achieve its own policy goals. This example expands on the main example from Kruse and Strack [24] by adding the possibility that the worker can choose to leave the market.

- (Hypothesis Testing) The principal hires an econometrician to conduct a hypothesis test. The econometrician performs a Bayesian sequential test of H_0 vs H_1. After obtaining each sample, she can choose to reject one of the two hypotheses or obtain an additional sample. Payments are based on sample size to incentivize the econometrician to perform the designed test.

- (Project Funding) The principal, for instance, a public authority, a government agency sponsoring research, or a city council, is deciding which one of two projects it will invest extra resources into and wants to hire an investigator for advice. The investment decision affects the utility of the investigator, so the principal would like to find a payment scheme to incentivize the investigator to report truthfully.

In the above examples, the strategies for the agent consist of:

- (A stopping rule) The stopping rule can be complicated, but it often involves a sequence of pairs of thresholds that may vary over time. If the state falls inside the thresholds, the agent will continue; otherwise, she stops.

- (A choice of alternative) After the agent stops, she makes a choice between the two alternatives.

The above class of strategies is defined formally as the two-sided threshold rules. I analyze the implementability of this class of strategies.

The main result describes the necessary and sufficient conditions for the intrinsic utility and Markov process that ensure the implementability of all two-sided threshold rules. The conditions are similar to the single crossing condition in Kruse and Strack [24] and the monotonicity condition from Pavan, Segal, and Toikka [33]. In addition, closed form formulas for the transfers that implement these stopping rules can be found.

The agent's problem has been studied since Wald (1947) and Arrow, Blackwell, and Girshick (1949) as the optimal stopping problem. The principal's problem, which is the focus of this analysis, is most closely related to Kruse and Strack [24], where the principal chooses transfers to influence the stopping rule used by the agent. Kruse and Strack [24] define a cut-off stopping rule as a strategy where the agent stops the first time the state she privately observes is above a certain threshold. They make a single crossing assumption (KS-SCC) that is a sufficient and necessary condition for the implementability of cut-off rules. It requires the expected difference between utilities in consecutive periods, which they call marginal incentives, to be weakly decreasing in the state. They also give closed form solutions for the transfers.

The model is also related to Pavan, Segal, and Toikka [33]. In their paper, there are multiple agents, and each agent can choose an allocation after each period. They show that a sequence of allocations is implementable if and only if an integral monotonicity condition holds. They also have a stronger condition called the single crossing condition (PST-SCC), which is a sufficient condition for implementability. PST-SCC requires the expected sums of discounted marginal utilities to be monotonic in state. The discount factor they use is the impulse response function. They give a characterization of the transfers but not closed form formulas for the transfers.

Due to the assumption that the agent can choose one of two actions only once, not only am I able to simplify PST-SCC, but the conditions in this model will also hold for utility functions that are non-monotonic in a particular way that does not satisfy KS-SCC. If the utilities are increasing in state for one alternative and decreasing in state for the other, then these conditions imply implementability as well.

I show that the condition of implementability can be simplified for the previously discussed examples:

- (Job Search) All stopping rules with two-sided thresholds are implementable if the expected change in utility from choosing the same action in any two consecutive periods is weakly decreasing in the current state.

- (Hypothesis Testing) As long as the payoff from accepting the more likely hypothesis is higher, the Bayesian sequential test with any size is implementable.

- (Project Funding) Suppose the investigator observes a state process that is additive, the payoff is linear in the states, and the slopes have different signs for different projects. Then, every strategy where the investigator stops when the state is either too high or too low is implementable.

Section 2 introduces the model, Section 3 gives sufficient and necessary conditions when threshold rules are implementable, and Section 4 provides simplifications of these conditions for examples with special utilities or stochastic processes of the states.

3.2 Model

3.2.1 Agent's Problem

Consider a stopping problem with two possible terminal actions. In each period $t = 0, 1, ..., T$, an agent observes the state of a one-dimensional Markov process, $x_t \in [\underline{x}, \bar{x}]$. Next, she chooses between three options: stop and choose action -1, stop and choose action +1, or continue. If the agent stops at time $\tau = t$ and chooses $q_t \in \{-1, +1\}$, she receives utility, $u_t(q_t, x_t)$.

I make the following monotonicity assumption on the utility function. It is a variation of a part of the Spence-Mirrlees conditions that is necessary for implementability.

Assumption 1. The period t utility functions are partially differentiable with respect to x for all $t \in \{0, 1, ..., T\}$, and,

$$\frac{\partial u_t(-1, x)}{\partial x} \leq 0 \text{ and } \frac{\partial u_t(+1, x)}{\partial x} \geq 0, \ \forall\, x \in [\underline{x}, \bar{x}].$$

A strategy (or decision rule) is a sequence of mappings $q_t : [\underline{x}, \bar{x}] \to Q$ where $Q = \{-1, 0, 1\}$ is the set of decisions. The decision q_t is independent of previous states $x_0, x_1, ..., x_{t-1}$ because the utility only depends on the current x_t the agent observes. Each strategy induces a stopping time, which is a mapping $\tau : [\underline{x}, \bar{x}]^T \to \{0, 1, ..., T\}$ defined so that,

$$\tau(x_0, ..., x_T) = \min\{t : q_t(x_t) \neq 0\}.$$

3.2.2 Implementation Problem

The principal chooses transfers, $p_t(-1)$ and $p_t(+1)$, in order to provide the agents with incentives to pick a particular strategy (or decision rule), q_t. As in Kruse and Strack [24], I only consider posted price mechanisms. The prices only depend on the stopping time and the action the agent chooses, but not the exact history of reports. The agent observing state x_t in period t gets payoff $u_t(q_t, x_t) - p_t(q_t)$ if she stops and makes a decision, $q_t \in \{-1, +1\}$.

The expected utility at time 0 for the agent is,

$$U_0(q, x) = \mathbb{E}\left[u_\tau(q_\tau(X_\tau), X_\tau) - p_\tau(q_\tau(X_\tau)) \,|X_0 = x\right]$$

where τ is the stopping time implied by the strategy q.

Definition 16. A strategy, $q = \{q_t\}_{t=0}^T$, is implementable if there are transfers $\{p_t\}_{t=0}^T$ such that for an agent observing state $x \in [\underline{x}, \bar{x}]$ at time $t = 0$,

$$\{q_t\}_{t=0}^T = \arg\max_{\{q'_t\}_{t=0}^T} \mathbb{E}\left[u_{\tau'}(q'_{\tau'}(X_{\tau'}), X_{\tau'}) - p_{\tau'}(q'_{\tau'}(X_{\tau'})) \,|X_0 = x\right].$$

A strategy (with its implied stopping time) is implementable if it is possible to provide the agent with incentives to choose that strategy.

3.2.3 Threshold Rules

I am interested in threshold strategies where the agent stops whenever the state falls outside of a sequence of intervals, $\{[a_t, b_t]\}_{t=0}^{T}$.

Definition 17. A strategy, q, is a (two-sided) threshold strategy if there are sequences $\{a_t, b_t\}_{t=0}^{T}$ such that $a_t \le b_t, a_T = b_T$ and,

$$q_t = \begin{cases} -1 & \text{if } x_t \in [\underline{x}, a_t) \\ 0 & \text{if } x_t \in [a_t, b_t] \\ +1 & \text{if } x_t \in (b_t, \bar{x}] \, . \end{cases}$$

I denote a threshold strategy with thresholds $\{a_t, b_t\}_{t=0}^{T}$ as $q^{(a,b)}$. Kruse and Strack [24] uses cut-off stopping rules with only one sequence of thresholds, and such rules are special cases where $a_t = \underline{x}$.

As presented in the next section, all two-sided threshold strategies are implementable, but it is not true that all implementable strategies have a stopping rule in the form of two-sided thresholds. For example, with non-concave utility functions, strategies in which the agent stops when the state falls outside multiple intervals in the same period can be implementable. Chow, Robbins, and Siegmund [11] provides examples of such stochastic processes when the optimal stopping rule is not a two-sided threshold strategy.

3.2.4 Stochastic Process of the States

I follow assumptions similar to Kruse and Strack [24] on the stochastic process, $\{X_t\}_{t=0}^{T}$. These assumptions are needed for the model to be tractable.

Assumption 2. (Regular Transitions) The process $\{X_t\}_{t=0}^{T}$ satisfies,

1. For any continuous $\phi : [\underline{x}, \bar{x}] \to \mathbb{R}, \mathbb{E}\left[\phi\left(X_{t+1}\right)|X_t = x\right]$ is continuous in x.

2. For any weakly decreasing $\phi : [\underline{x}, \bar{x}] \to \mathbb{R}, \mathbb{E}\left[\phi\left(X_{t+1}\right)|X_t = x\right]$ is non-increasing in x.

3. For any interval $[a, b) \subseteq [\underline{x}, \bar{x}]\,, \mathbb{E}\left[1_{X_{t+1} \in [a,b)}|X_t = x\right] > 0$ for each $x \in [\underline{x}, \bar{x}]$.

The continuity-preserving and monotonicity-preserving properties are used to ensure that the shape of the value function in the future periods stays the same after taken expectations given the state of the current period. The full support property is included to ensure the uniqueness of the representation of a threshold rule. Without full support, multiple thresholds can be used to represent the same stopping time.

3.3 Implementability

Theorem 1 finds conditions under which all threshold rules are implementable.

Let F_t be the conditional cumulative distribution function of X_t given X_{t-1}, and f_t be the conditional density function of $X_t|X_{t-1}$ for the Markov process.

Definition 18. Define the impulse response function as,

$$\mathcal{I}\left(x_{t+1}, x_t\right) = -\frac{\partial F_{t+1}\left(x_{t+1}|x_t\right)}{\partial x_t} \frac{1}{f_{t+1}\left(x_{t+1}|x_t\right)}.$$

The impulse response function is used in Pavan, Segal, and Toikka [33] to state their monotonicity conditions, too. It is used as a discount factor for marginal utilities so that they can be added over time.

Theorem 2. *The following are equivalent:*

1. *(Implementability) Every two-sided threshold rule, $q^{(a,b)}$, is implementable.*

2. *(Monotonic Marginal Incentive Condition) The following marginal incentive function is weakly decreasing in x for each $t \in \{0, 1, ..., T\}$.*

$$\mathbb{E}\left[u_{t+1}\left(s, X_{t+1}\right)|X_t = x\right] - u_t\left(s, x\right), \text{ for } s \in \{-1, +1\}$$

3. *(Single Crossing Condition) The following single crossing condition is satisfied for each $t \in \{0, 1, ..., T\}$ and $x \in [\underline{x}, \bar{x}]$.*

$$s \cdot \mathbb{E}\left[\frac{\partial u_{t+1}\left(s, X_{t+1}\right)}{\partial x}\mathcal{I}\left(X_{t+1}, X_t\right)|X_t = x\right] \leq s \cdot \frac{\partial u_t\left(s, x\right)}{\partial x}, \text{ for } s \in \{-1, +1\}$$

Both the marginal incentive condition and the single crossing condition are sufficient to implement any threshold rule but are not necessary to implement a particular threshold rule. They are only necessary to implement all possible threshold rules at the same time.

The two conditions for implementability are equivalent due to a simple integration by parts. The sufficiency of the single crossing condition (3) $\Rightarrow$ (1) can be proven using results in Pavan, Segal, and Toikka [33]. In particular, the assumptions about the utility function and the single crossing conditions imply strong monotonicity, which is a sufficient condition for implementation from their paper. I present an alternative proof for this specific model, which, in addition to proving implementability, also gives a closed form formula for the transfers and contains parts that are useful for the proof of the necessity of the single crossing conditions in implementing threshold rules. The sufficiency of the marginal incentives (2) $\Rightarrow$ (1) for one-sided thresholds with $a_t = \underline{x}$ is shown in Kruse and Strack [24], and I modify and extend their proof for two-sided thresholds, where the principal needs to prevent the agent from choosing both actions -1 and +1 in future periods if he wants the agent to stop in the current period.

The proof for the necessity of the two conditions for implementing all threshold rules (1) $\Rightarrow$ (2) or (3) is new and does not rely on the techniques or results from Kruse and Strack [24] or Pavan, Segal, and Toikka [33].

In order to describe these transfers, additional notations are needed.

Definition 19. Define the modified marginal incentive functions as,

$$z_t(x) = \mathbb{E}\left[\max_{s'}\left\{u_{t+1}(s', X_{t+1})\right\}|X_t = x\right] - \max_{s'}\left\{u_t(s', x)\right\},$$
$$\tilde{z}_t(s, x) = \mathbb{E}\left[\max_{s'}\left\{u_{t+1}(s', X_{t+1})\right\}|X_t = x\right] - u_t(s, x), \text{ for } s \in \{+1, -1\}.$$

The function $z_t(x)$ is the expected gain in utility if the agent continues for another period relative to stopping in period t. The function $\tilde{z}_t(s, x)$ is the expected gain in utility if the agent continues for another period relative to reporting the state she is supposed to.

Corollary 3. *The following sequence of transfers implement $q^{(a,b)}$.*

$$p_t(-1) = \tilde{z}_t(-1, a_t) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s\left(\min\left\{\max\left\{a_t, X_s\right\}, b_t\right\}\right)|X_t = a_t\right]$$
$$p_t(+1) = \tilde{z}_t(+1, b_t) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s\left(\min\left\{\max\left\{a_t, X_s\right\}, b_t\right\}\right)|X_t = b_t\right]$$

The expressions are similar to the ones given in Kruse and Strack [24], but the formulas for both z_t and $\tilde{z}_t$ are different from the one defined in their paper and are also different from the expression in the marginal incentive condition in Theorem 2. The monotonicity of the price functions is the key part in proving (2) and (3) $\Rightarrow$ (1) in Theorem 2.

Throughout the paper, I assumed that the principal has an optimal stopping rule he plans to implement. The paper does not mention how such rule is chosen, however, given an objective function, for example, profit maximization, it is possible to find the optimal sequence of thresholds since the expressions of the transfers in terms of the thresholds are provided for every two-sided threshold strategy.

3.4 Examples

Example 3. A worker observes job offers with wage $x_t \in [0, \bar{x}]$ in period t and decides whether to accept the offer, reject the offer and permanently leave the market, or keep searching. The government pays lump sum employment insurance , p_τ, in period τ to incentivize the worker to accept offers above some upper threshold and to exit the market when offers are below some lower threshold. Alternatively, the government can use other tax or transfer schemes to influence the worker into accepting offers earlier or to stop searching and get additional education. This example extends the one in Kruse and Strack [24] to include the possibility of the worker exiting the market.

For example, if the utility the worker gets is equal to the wage when she accepts the offer, and 0 when she stops searching permanently, the conditions in Theorem 2 become,

$$z_t\left(x\right) = \mathbb{E}\left[X_{t+1}|X_t = x\right] - x \text{ is weakly decreasing in } x.$$

If the above condition holds, then the following transfers can be used to incentivize the worker to choose a job search strategy described by any thresholds $\{(a_t, b_t)\}_{t=0}^{T.}$

$$p_t\left(\text{ exit market }\right) = \mathbb{E}\left[X_{t+1}|X_t = a_t\right] + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s\left(\min\left\{\max\left\{a_t, X_s\right\}, b_t\right\}\right)|X_t = a_t\right]$$

$$p_t\left(\text{ accept offer }\right) = z_t\left(b_t\right) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s\left(\min\left\{\max\left\{a_t, X_s\right\}, b_t\right\}\right)|X_t = b_t\right]$$

The transfers are different from the expressions in Kruse and Strack [24] due to the addition of lower thresholds $\{a_t\}_{t=0}^{T}$. If $a_t = 0$ for $t = 0, 1, 2, ..., T$, then the condition and transfer functions coincide with the ones in their paper.

Example 4. A Bayesian statistician must distinguish between the two possible states of the world, $s = -1$ or $s = +1$. The statistician starts with a belief, x_0, that the state is $+1$. The statistician observes data and uses it to update the belief to posterior x_t. In each period t, she can either decide to collect more data at cost c_t, or choose one of two possible states. If in period t, she decides that the state of the world is $q_t = s$ and the true state of the world is also s, her payoff is equal to $\alpha_t\left(s\right)$. Otherwise, if the true state is -s, she gets $\beta_t\left(s\right)$, for each t and $s \in \{-1, +1\}$. Define the following constants,

$$m_t\left(-1\right) = \alpha_t\left(-1\right) - \beta_t\left(-1\right) \text{ and } m_t\left(+1\right) = \beta_t\left(+1\right) - \alpha_t\left(+1\right),$$

$$k_t\left(-1\right) = \beta_t\left(-1\right) - \sum_{i=0}^{t} c_i \text{ and } k_t\left(+1\right) = \alpha_t\left(+1\right) - \sum_{i=0}^{t} c_i.$$

Then, the utility function is linear in the posterior belief,

$$u_t\left(q_t, x_t\right) = m_t\left(q_t\right) x_t + k_t\left(q_t\right),$$

and Assumption 1 is equivalent to $\beta_t\left(s\right) < \alpha_t\left(s\right)$ for each t.

The single crossing condition in Theorem 2 is simplified to,

$$\mathbb{E}\left[\mathcal{I}\left(X_{t+1}, X_t\right) | X_t = x\right] \leq \min\left\{\frac{m_t\left(+1\right)}{m_{t+1}\left(+1\right)}, \frac{m_t\left(-1\right)}{m_{t+1}\left(-1\right)}\right\}.$$

If the belief process forms a martingale and $\alpha_t\left(s\right) = \alpha\left(s\right), \beta_t\left(s\right) = \beta\left(s\right)$ for each t, then the single crossing condition in Theorem 2 is always satisfied due to the following derivation,

$$\mathbb{E}\left[\mathcal{I}\left(X_{t+1}, X_t\right) | X_t = x\right] = \frac{d}{dx}\left(\mathbb{E}\left[X_{t+1} | X_t = x\right]\right) = 1,$$

because of the martingale property, and since α and β are constant over time.

$$\frac{m_{t+1}\left(-1\right)}{m_t\left(-1\right)} = \frac{m_{t+1}\left(+1\right)}{m_t\left(+1\right)} = 1$$

Therefore, the single crossing condition is always satisfied. Bayesian hypothesis test of any size can be implemented by adding the $\{p_t\}_{t=0}^{T}$ from Theorem 2 to the loss function or, equivalently, to the cost of the samples.

Example 5. A government agency sponsoring research is choosing which one of two universities to invest in and hires an investigator for advice. The investigator observes an additive state process, $X_{t+1} = X_t + \varepsilon_t$, where $\varepsilon_t \sim G_t$, for some independently distributed G_t and the utility is linear in the state with a constant discount factor $\delta < 1$.

$$u_t\left(q_t, x_t\right) = \delta^t\left(m_t\left(q_t\right) x_t + k_t\left(q_t\right)\right)$$

Then, the single crossing condition in Theorem 2 is always satisfied due to the following derivation,

$$\mathcal{I}\left(x_{t+1}, x\right) = -\left(\frac{\partial G_t\left(x_{t+1} - x\right)}{\partial x}\right)\frac{1}{g_t\left(x\right)} = 1,$$

because of the additivity of the process, and,

$$\frac{m_t(-1)}{m_{t+1}(-1)} = \frac{m_t(+1)}{m_{t+1}(+1)} = \frac{1}{\delta}.$$

Therefore, linear utility with either martingale processes or additive processes leads to implementability.

3.5 Proof of Theorem 1

Define the expression in Theorem 2 as $\tilde{\tilde{z}}$.

$$\tilde{\tilde{z}}_t(s,x) = \mathbb{E}\left[u_{t+1}(s, X_{t+1})\,|X_t = x\right] - u_t(s,x)\,, \text{ for } s \in \{-1,+1\}$$

This is the expected gain in utility for waiting for one more period while being forced to report the same state, s.

Equivalence (2) $\Leftrightarrow$ (3): I start by converting the single crossing condition in Theorem 2 to the monotonicity condition of $\tilde{\tilde{z}}_t$.

Lemma 15. *The single crossing condition in Theorem 2 holds if and only if $s \cdot \tilde{\tilde{z}}_t(s,x)$ is weakly decreasing for $s \in \{-1,+1\}$.*

Implication (2) $\Rightarrow$ (1): I prove the monotonicity of $s \cdot \tilde{\tilde{z}}_t(s,x)$ implies implementability, by starting with the following observation on $\tilde{z}_t$.

Lemma 16. *If $s \cdot \tilde{\tilde{z}}_t(s,x)$ is weakly decreasing, then $s \cdot \tilde{z}_t(s,x)$ is always weakly decreasing.*

Define the following transfer functions. They are not the actual transfers because they also depend on x.

$$p_t(s,x) = \tilde{z}_t(s,x) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s\left(\min\left\{\max\left\{a_t, X_s\right\}, b_t\right\}\right)|X_t = x\right]$$

The following lemma shows the monotonicity of $p_t(s,x)$.

Lemma 17. *If $s \cdot \tilde{\tilde{z}}_t(s,x)$ is weakly decreasing, then $s \cdot p_t(s,x)$ is weakly decreasing.*

The value function has the following form by induction.

Lemma 18. *The value function is,*

$$V_t(x) = u_t(x) + \sum_{s=t}^{T-1} \mathbb{E}\left[z_s\left(\min\left\{\max\left\{a_s, X_s\right\}, b_s\right\}\right)|X_t = \min\left\{\max\left\{a_t, x\right\}, b_t\right\}\right].$$

At the end, I show that, given the above value function, any threshold strategy is implementable.

Lemma 19. *Given the value function from Lemma 18, $q^{(a,b)}$ is implementable.*

Implication (1) $\Rightarrow$ (2): I prove implementability implies the monotonicity of $s \cdot \tilde{\tilde{z}}_t(s,x)$.

Lemma 20. *Suppose $q^{(a,b)}$ is implementable for any thresholds (a,b), then $s \cdot \tilde{\tilde{z}}_t(s,x)$ is weakly decreasing.*

The following are the proofs of the lemmas used.

Proof of Lemma 15:

$$
\begin{aligned}
s \cdot \frac{d\tilde{\tilde{z}}_t(s,x)}{dx} &= s \cdot \frac{d}{dx}\mathbb{E}\left[u_{t+1}(s, X_{t+1}) \,|X_t\right] - s \cdot \frac{\partial u_t(s,x)}{\partial x} \\
&= -s \cdot \int_{\underline{x}}^{\bar{x}} \frac{\partial u_{t+1}(s, x_{t+1})}{\partial x} \frac{\partial F_{t+1}(x_{t+1}|x)}{\partial x} dx_{t+1} - s \cdot \frac{\partial u_t(s,x)}{\partial x} \\
&= s \cdot \int_{\underline{x}}^{\bar{x}} \frac{\partial u_{t+1}(s, x_{t+1})}{\partial x} \mathcal{I}(X_{t+1}, x) f_{t+1}(x_{t+1}|x)\, dx_{t+1} - s \cdot \frac{\partial u_t(s,x)}{\partial x} \\
&= s \cdot \mathbb{E}\left[\frac{\partial u_{t+1}(s, X_{t+1})}{\partial x} \mathcal{I}(X_{t+1}, X_t) \,|X_t = x\right] - s \cdot \frac{\partial u_t(s,x)}{\partial x}
\end{aligned}
$$

Therefore, $s \cdot \tilde{\tilde{z}}_t(s,x)$ is weakly decreasing if and only if,

$$
s \cdot \mathbb{E}\left[\frac{\partial u_{t+1}(s, X_{t+1})}{\partial x} \mathcal{I}(X_{t+1}, X_t) \,|X_t = x\right] \leq s \cdot \frac{\partial u_t(s,x)}{\partial x},
$$

which is the condition in Theorem 2. □

Proof of Lemma 16: To shorten the notations, define,

$$
u_t(x) = \max_{s'} \left\{u_{t+1}(s', X_{t+1})\right\}.
$$

Due to Assumption 1, following inequality holds,

$$
\frac{du_t(-1,x)}{dx} \leq \frac{du_t(x)}{dx} \leq \frac{du_t(+1,x)}{dx}.
$$

Therefore, $s \cdot \tilde{z}_t(s,x)$ is also weakly decreasing if the condition in Theorem 2 holds.

$$
s \cdot \frac{d\tilde{z}_t(s,x)}{dx} \leq s \cdot \frac{d\tilde{\tilde{z}}_t(s,x)}{dx} \leq 0
$$

□

Proof of Lemma 17: When $t = T - 1$,

$$s \cdot p_t(s, x) = s \cdot \tilde{z}_{T-1}(s, x),$$

which is weakly decreasing from Lemma 16.

The rest of the proof goes by backward induction on t. Assuming $s \cdot p_{t+1}(s, x)$ is weakly decreasing,

$$\begin{aligned}
s \cdot p_t(s, x) &= s \cdot \left(\tilde{z}_t(s, x) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\}) \,|X_t = x \right] \right) \\
&= s \cdot \left(\tilde{z}_t(s, x) + \sum_{s=t+2}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\}) \,|X_t = x \right] \right) \\
&\quad + s \cdot \mathbb{E}\left[z_{t+1}(\min\{\max\{a_{t+1}, X_{t+1}\}, b_{t+1}\}) \,|X_t = x \right] \\
&= s \cdot \mathbb{E}\left[p_{t+1}(s, X_{t+1}) \,|X_t = x \right] \\
&\quad - s \cdot \mathbb{E}\left[u_{t+1}(s, \min\{\max\{a_t, X_{t+1}\}, b_t\}) \,|X_t = x \right] - s \cdot u_t(s, x) \\
&= \mathbb{E}\left[s \cdot p_{t+1}(s, X_{t+1}) \,|X_t = x \right] + s \cdot \tilde{\tilde{z}}_t(s, x).
\end{aligned}$$

Here, $s \cdot p_{t+1}(s, X_{t+1})$ is weakly decreasing by induction hypothesis, $s \cdot \tilde{\tilde{z}}_t(s, x)$ is weakly decreasing by assumption, and taking conditional expectations on X_t preserves monotonicity by Assumption 2. Therefore, $s \cdot p_t(s, x)$ is weakly decreasing in x for $s \in \{-1, +1\}$. □

Proof of Lemma 18: The base case when $t = T$ is,

$$V_T(x) = u_T(x) \text{ with } p_t(-1, a_t) = p_t(+1, b_t) = 0.$$

If $p_T \neq 0$ at $a_T = b_T$, redefine $u_t(s, x)$ by subtracting the $p_t(-1, a_t)$ from it for all t.

For the induction, I assume that,

$$V_{t+1}(x) = u_{t+1}(x) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\}) \,|X_{t+1} = \min\{\max\{a_t, x\}, b_t\} \right].$$

For $x \leq x_t^\star$, the value function can be simplified to,

$$\begin{aligned}
V_t(x) &= \max\left\{u_t(-1,x) - p_t(-1,a_t), u_t(1,x) - p_t(+1,b_t), \mathbb{E}\left[V_{t+1}(X_{t+1})|X_t = x\right]\right\} \\
&= \max\left\{u_t(-1,x) - p_t(-1,a_t), \mathbb{E}\left[V_{t+1}(X_{t+1})|X_t = x\right]\right\} \\
&= u_t(-1,x) + \max\left\{\tilde{z}_t(-1,a_t) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\left\{\max\left\{a_t, X_s\right\}, b_t\right\})|X_t = a_t\right],\right. \\
&\quad \left.\tilde{z}_t(-1,x) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\left\{\max\left\{a_t, X_s\right\}, b_t\right\})|X_t = x\right]\right\} \\
&= u_t(-1,x) + \tilde{z}_t(-1, \max\left\{a_t, x\right\}) \\
&\quad + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\left\{\max\left\{a_s, X_s\right\}, b_s\right\})|X_t = \max\left\{a_t, x\right\}\right] \\
&= u_t(-1,x) + \mathbb{E}\left[u_{t+1}(X_{t+1})|X_t = \max\left\{a_t, x\right\}\right] - u_t(-1, \max\left\{a_t, x\right\}) \\
&\quad + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\left\{\max\left\{a_s, X_s\right\}, b_s\right\})|X_t = \max\left\{a_t, x\right\}\right].
\end{aligned}$$

The second to last line is obtained by substituting in the transfers and applying the induction hypothesis.

□

Proof of Lemma 19: I define $U_t(x)$ as the utility the agent observing state x gets if she stops.

$$U_t(x) = \max\left\{u_t(-1,x) - p_t(-1,a_t), u_t(+1,x) - p_t(+1,b_t)\right\}$$

Let $x_t^\star$ be the x that satisfies $u_t(-1,x) - p_t(-1,a_t) = u_t(+1,x) - p_t(+1,b_t)$. Then due to the monotonicity of $u_t(s,x)$ from the assumption of the lemma, $U_t(x)$ can be rewritten as,

$$U_t(x) = \begin{cases} u_t(-1,x) - p_t(-1,a_t) & \text{if } x \leq x_t^\star \\ u_t(+1,x) - p_t(+1,b_t) & \text{if } x \geq x_t^\star. \end{cases}$$

If $x > x_t^\star$, using similar arguments, we obtain,

$$\begin{aligned}
V_t(x) = u_t(+1,x) &+ \mathbb{E}\left[u_{t+1}(X_{t+1})|X_t = \min\left\{x, b_t\right\}\right] - u_t(+1, \min\left\{x, b_t\right\}) \\
&+ \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\left\{\max\left\{a_s, X_s\right\}, b_s\right\})|X_t = \max\left\{a_t, x\right\}\right].
\end{aligned}$$

Combining the above two pieces results in the desired form,

$$\begin{aligned}
V_t(x) = u_t(x) + \mathbb{E}\left[u_{t+1}(X_{t+1}) \,|X_t = \min\{\max\{a_t, x\}, b_t\}\right] - u_t(\min\{\max\{a_t, x\}, b_t\}) \\
+ \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\}) \,|X_t = \min\{\max\{a_t, x\}, b_t\}\right] \\
= u_t(x) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\})\right] + z_t(\min\{\max\{a_t, x\}, b_t\}) \\
= u_t(x) + \sum_{s=t}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\})\right].
\end{aligned}$$

Therefore, if $x < a_t$,

$$\begin{aligned}
V_t(x) - U_t(x) &= V_t(x) - (u_t(-1, x) - p_t(-1, a_t)) \\
&= \tilde{z}_t(-1, a_t) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\}) \,|X_t = a_t\right] \\
&\quad - \tilde{z}_t(-1, a_t) - \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\}) \,|X_t = a_t\right] \\
&= 0.
\end{aligned}$$

The same expression can be found for $x > b_t$.

If $a_t \le x \le x_t^\star$,

$$\begin{aligned}
V_t(x) - U_t(x) &= V_t(x) - (u_t(-1, x) - p_t(-1, a_t)) \\
&= \tilde{z}_t(-1, x) + \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\}) \,|X_t = x\right] \\
&\quad - \tilde{z}_t(-1, a_t) - \sum_{s=t+1}^{T-1} \mathbb{E}\left[z_s(\min\{\max\{a_s, X_s\}, b_s\}) \,|X_t = a_t\right] \\
&\ge 0.
\end{aligned}$$

Similarly, if $x_t^\star \le x \le b_t$,

$$V_t(x) - U_t(x) \ge 0,$$

and if $x \geq b_t$,

$$V_t(x) - U_t(x) = 0.$$

Combining the four pieces for x in different parts of the domain, we get,

$$V_t(x) - U_t(x) \begin{cases} = 0 & \text{if } x \notin [a_t, b_t] \\ \geq 0 & \text{if } x \in [a_t, b_t]. \end{cases}$$

Therefore, the optimal strategy is $q^{(a,b)}$, and this is implementable by transfers $p_t(-1, a_t)$ and $p_t(+1, b_t)$, which are the transfers given in Lemma 17. □

Proof of Lemma 20: Consider the two sequences of thresholds $\{(a_t, \bar{x})\}_{t=0}^{T}$ and $\{(\underline{x}, b_t)\}_{t=0}^{T}$.

Given $\{(\underline{x}, b_t)\}_{t=0}^{T}$ is implementable, take $x < x'$ and assume $\tilde{\tilde{z}}_t(+1, x) < \tilde{\tilde{z}}_t(+1, x')$ for a contradiction.

For x, the value function in the recursive form is given by,

$$\begin{aligned} V_t(x) - U_t(x) &= \max\{U_t(x), \mathbb{E}[V_{t+1}(X_{t+1}) | X_t = x]\} - U_t(x) \\ &= \max\{0, \mathbb{E}[V_{t+1}(X_{t+1}) - u_{t+1}(X_{t+1}) | X_t = x] \\ &\quad + \mathbb{E}[u_{t+1}(X_{t+1}) | X_t = x] - u_t(+1, x) - p_t(-1, a_t)\} \\ &= \max\left\{0, \mathbb{E}[V_{t+1}(X_{t+1}) - u_{t+1}(X_{t+1}) | X_t = x] + \tilde{\tilde{z}}_t(+1, x) - p_t(-1, a_t)\right\}. \end{aligned}$$

All of the future transfers are fixed for different values of b_t. Therefore, $V_{t+1}(x)$ and $U_t(x)$ are fixed for different values of b_t. Consider two cases for b_t as follows:

For $x < x'$, if the following holds,

$$\mathbb{E}[V_{t+1}(X_{t+1}) - u_{t+1}(X_{t+1}) | X_t = x] \leq \mathbb{E}[V_{t+1}(X_{t+1}) - u_{t+1}(X_{t+1}) | X_t = x'],$$

I implement the stopping rule $b_t = x'$. This means that the agent observing state $x_t = x$ should continue at time t. However,

$$\begin{aligned} &\mathbb{E}[V_{t+1}(X_{t+1}) - u_{t+1}(X_{t+1}) | X_t = x] + \tilde{\tilde{z}}_t(+1, x) - p_t(-1, a_t) \\ &< \mathbb{E}[V_{t+1}(X_{t+1}) - u_{t+1}(X_{t+1}) | X_t = x'] + \tilde{\tilde{z}}_t(+1, x') - p_t(-1, a_t) \\ &= 0. \end{aligned}$$

The inequality shows that it is optimal for the agent observing state $x_t = x$ to stop, contradicting the definition of the stopping rule with $b_t = x' > x$.

Similarly, if the following holds,

$$\mathbb{E}\left[V_{t+1}\left(X_{t+1}\right) - u_{t+1}\left(X_{t+1}\right)|X_t = x\right] \geq \mathbb{E}\left[V_{t+1}\left(X_{t+1}\right) - u_{t+1}\left(X_{t+1}\right)|X_t = x'\right],$$

I implement the stopping rule $b_t = x$. It means that the agent observing state $x_t = x$ should stop at time t. However,

$$\begin{aligned}
&\mathbb{E}\left[V_{t+1}\left(X_{t+1}\right) - u_{t+1}\left(X_{t+1}\right)|X_t = x'\right] + \tilde{\tilde{z}}_t\left(+1, x'\right) - p_t\left(-1, a_t\right) \\
&> \mathbb{E}\left[V_{t+1}\left(X_{t+1}\right) - u_{t+1}\left(X_{t+1}\right)|X_t = x\right] + \tilde{\tilde{z}}_t\left(+1, x\right) - p_t\left(-1, a_t\right) \\
&= 0.
\end{aligned}$$

This shows that it is optimal for the agent observing state $x_t = x'$ to continue, contradicting the definition of the stopping rule with $b_t = x < x'$.

Therefore, $\tilde{\tilde{z}}_t\left(+1, x\right)$ is weakly decreasing in x.

Similarly, $\tilde{\tilde{z}}_t\left(-1, x\right)$ is weakly increasing in x since $\{(a_t, \bar{x})\}_{t=0}^{T}$ is implementable.

Therefore, $s \cdot \tilde{\tilde{z}}_t\left(s, x\right)$ is weakly decreasing in x for $s \in \{+1, -1\}$. $\square$

Bibliography

[1] Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda (2011), "Resolving conflicting preferences in school choice: The 'boston mechanism' reconsidered." *American Economic Review*, 101, 399–410.

[2] Abdulkadiroğlu, Atila, Parag A Pathak, and Alvin E Roth (2009), "Strategy-proofness versus efficiency in matching with indifferences: Redesigning the nyc high school match." *American Economic Review*, 99, 1954–78.

[3] Albrecht, James, Axel Anderson, and Susan Vroman (2010), "Search by committee." *Journal of Economic Theory*, 145, 1386–1407.

[4] Athey, Susan and Ilya Segal (2013), "An efficient dynamic mechanism." *Econometrica*, 81, 2463–2485.

[5] Bergemann, Dirk and Stephen Morris (2005), "Robust mechanism design." *Econometrica*, 73, 1771–1813.

[6] Bergemann, Dirk and Juuso Välimäki (2010), "The dynamic pivot mechanism." *Econometrica*, 78, 771–789.

[7] Board, Simon and Andrzej Skrzypacz (2016), "Revenue management with forward-looking buyers." *Journal of Political Economy*, 124, 1046–1087.

[8] Bogomolnaia, Anna and Hervé Moulin (2001), "A new solution to the random assignment problem." *Journal of Economic Theory*, 100, 295–328.

[9] Che, Yeon-Koo and Fuhito Kojima (2010), "Asymptotic equivalence of probabilistic serial and random priority mechanisms." *Econometrica*, 78, 1625–1672.

[10] Che, Yeon-Koo and Konrad Mierendorff (2016), "Optimal sequential decision with limited attention." *Unpublished manuscript, Columbia University.*

[11] Chow, Yuan Shih, Herbert Robbins, and David Siegmund (1991), *The Theory of Optimal Stopping.* Dover.

[12] Compte, Olivier and Philippe Jehiel (2010), "Bargaining and majority rules: A collective search perspective." *Journal of Political Economy*, 118, 189–221.

[13] Eső, Péter and Balazs Szentes (2007), "Optimal information disclosure in auctions and the handicap auction." *Review of Economic Studies*, 74, 705–731.

[14] Gershkov, Alex, Benny Moldovanu, and Xianwen Shi (2017), "Optimal voting rules." *Review of Economic Studies*, 84, 688–717.

[15] Gershkov, Alex, Benny Moldovanu, Xianwen Shi, et al. (2013), "Optimal mechanism design without money." Technical report.

[16] Gershkov, Alex, Benny Moldovanu, and Philipp Strack (2014), "Dynamic allocation and learning with strategic arrivals." *Unpublished manuscript, Hebrew University of Jerusalem.*

[17] Guo, Yingni (2016), "Dynamic delegation of experimentation." *The American Economic Review*, 106, 1969–2008.

[18] Guo, Yingni and Johannes Hörner (2015), "Dynamic mechanisms without money." Technical report, Reihe Ökonomie/Economics Series, Institut für Höhere Studien (IHS).

[19] Jarman, Felix, Vincent Meisner, et al. (2017), "Deterministic mechanisms, the revelation principle, and ex-post constraints." Technical report, CRC TRR 190 Rationality and Competition.

[20] Johnson, Terence R (2014), "Dynamic mechanism design without transfers: Promises and confidentiality." Technical report.

[21] Kamada, Yuichiro and Nozomu Muto (2011), "Multi-agent search with deadline." Technical report, mimeo.

[22] Kojima, Fuhito and Mihai Manea (2010), "Incentives in the probabilistic serial mechanism." *Journal of Economic Theory*, 145, 106–123.

[23] Kovác, Eugen, Daniel Krähmer, and Tymon Tatur (2013), "Optimal stopping in a principal-agent model with hidden information and no monetary transfers." *Unpublished manuscript.*

[24] Kruse, Thomas and Philipp Strack (2015), "Optimal stopping with private information." *Journal of Economic Theory*, 159, 702–727.

[25] Lipnowski, Elliot and Joao Ramos (2016), "Repeated delegation." *Unpublished manuscript.*

[26] Liu, Qingmin and Marek Pycia (2016), "Ordinal efficiency, fairness, and incentives in large markets." *Unpublished manuscript.*

[27] Moldovanu, Benny and Xianwen Shi (2013), "Specialization and partisanship in committee search." *Theoretical Economics*, 8, 751–774.

[28] Noda, Shunya (2016), "Full surplus extraction and within-period ex post implementation in dynamic environments." *Unpublished manuscript.*

[29] Parkes, David C, Ruggiero Cavallo, Florin Constantin, and Satinder Singh (2010), "Dynamic incentive mechanisms." *AI Magazine*, 31, 79–94.

[30] Parmigiani, Giovanni and Lurdes Inoue (2009), *Decision Theory: Principles and Approaches*, volume 812. John Wiley & Sons.

[31] Pathak, Parag A and Jay Sethuraman (2011), "Lotteries in student assignment: An equivalence result." *Theoretical Economics*, 6, 1–17.

[32] Pathak, Parag A and Tayfun Sönmez (2008), "Leveling the playing field: Sincere and sophisticated players in the boston mechanism." *American Economic Review*, 98, 1636–52.

[33] Pavan, Alessandro, Ilya Segal, and Juuso Toikka (2014), "Dynamic mechanism design: A myersonian approach." *Econometrica*, 82, 601–653.

[34] Peskir, Goran and Albert N Shiryaev (2006), *Optimal Stopping and Free-boundary Problems.* Springer Science & Business Media.

[35] Pycia, Marek (2014), "The cost of ordinality." *Unpublished manuscript.*

[36] Pycia, Marek and M Utku Ünver (2015), "Decomposing random mechanisms." *Journal of Mathematical Economics*, 61, 21–33.

[37] Strulovici, Bruno (2010), "Learning while voting: Determinants of collective experimentation." *Econometrica*, 78, 933–971.