University College London

Department of Economics

Applied Mechanism Design

Peter Postl

A thesis submitted for the degree of

Ph.D. in Economics

March 2004

Abstract

We use mechanism design techniques to analyze three novel problems. First, we study incentive contracts for information acquisition. In our setup, a principal must choose between two alternatives with unknown payoffs. An agent can find them out at a fixed cost per alternative. His effort and any information acquired are unobservable. The principal only observes the payoff of her chosen alternative, and hence the agent has the incentive to overstate cost by strategically finding out only one of the payoffs and lying about the other. We characterize sequential information acquisition procedures that can be delegated to the agent without granting him information rents.

Second, we study a situation where two agents have to choose one of three alternatives. Their ordinal rankings of these alternatives are diametrically opposed and are common knowledge. Ex ante efficiency requires that they implement the alternative that is ranked second by both if and only if the sum of their von Neumann Morgenstern utilities under this alternative is higher than under the two extreme options. Von Neumann Morgenstern utilities are privately observed types. We ask if there are incentive compatible mechanisms which elicit utilities and implement efficient decisions. We show that no such mechanisms exist if the distribution of agents' types has continuum support.

Finally, we investigate if procurement procedures that simultaneously determine specification and price of a good can result in an inefficient specification choice. In our setup, two suppliers can produce a good in either of two specifications which are equally good for the buyer. Costs are interdependent and unknown at the time of bidding. Each supplier receives a cost signal per specification. While an efficient mechanism exists, it involves a higher expected payment for the good than a mechanism that selects a bidder on the basis of price alone, in which case there is a chance of obtaining the specification with the highest production cost.

Contents

[1]This chapter is joint work with Tilman Börgers.

List of Figures

Acknowledgments

I am very grateful to my advisor Tilman Börgers for his guidance and support in the development of this thesis. Special thanks are also due to my second advisor Marco Ottaviani.

This thesis has also greatly benefitted from discussions with Gian Luigi Albano, Ralph Bailey, Ken Binmore, Jayasri Dutta, Antonella Ianni, Philippe Jehiel, Thomas Kittsteiner, Nathan Larson, Clare Leaver, Malcolm Pemberton and Nick Rau. I am grateful for their encouragement and support.

Financial support from the Economic and Social Research Council (ESRC), grant number R00429924359, is gratefully acknowledged.

Chapter 1

Introduction

Dispersion of relevant information is a pertinent feature of situations in which individuals strive to collectively reach a decision. Personal tastes, beliefs or experiences are often unknown to others, which makes communication of private information an integral part of the decision-making process. A mechanism is a collection of rules that state what form the individuals' communication may take and how it is aggregated into a collective decision. Examples of mechanisms abound: Auctions determine winner and selling price of an object on the basis of bids submitted by participants.[1] Insurance companies provide quotes on the basis of personal characteristics self-reported by potential clients.[2]

The purpose of Mechanism Design is to characterize mechanisms that give agents the incentive to communicate their private information truthfully.[3] We apply mechanism design techniques to three specific collective choice problems: We study contracts for information providers, the design of voting rules, and the construction of procurement procedures that elicit information about available product specifications. Before outlining these topics further, we present a brief overview of some important themes in

[1]See Krishna (2002).

[2]Mechanism Design has been used extensively in the contracting literature. For example, see Laffont and Martimort (2002) or Salanié (1998)).

[3]See Mas-Colell et al. (1995). Chapter 23.E provides a comprehensive overview of the Mechanism Design. Chapter 8.E describes games of incomplete information and the concept of Bayes Nash Equilibrium.

the mechanism design literature to which the present thesis aims to add a new angle.

At the heart of Mechanism Design is the question what properties mechanisms must have in order to induce truthful revelation of private information. In addressing this question, much attention in the literature has been devoted to collective decision problems (such as auctions) in which individuals' private information is captured by a scalar parameter (e.g. bidders' willingness to pay for an object at auction), utility is transferrable between individuals through money payments, and individuals' preferences are represented by linear utility functions.

In this setting, necessary and sufficient conditions for Bayesian incentive compatibility have been derived. For auctions, incentive compatibility implies that the expected utility of a bidder is determined up to a constant by the allocation rule that indicates with which probability the bidder obtains the object. A generalization of this result to settings in which individuals' private information has more than one dimension can be found in Jehiel and Moldovanu (2001) and Krishna and Perry (2000).

A second important theme in Mechanism Design is the question of pareto-optimality (or ex post efficiency) of collective decisions generated by incentive compatible mechanisms. If individuals in a (quasi-) linear collective choice setting can be compelled to participate, then there always exists at least one incentive compatible mechanism that achieves ex post efficiency.[4]

Myerson and Satterthwaite (1983) show that in bilateral trade settings in which the buyer and seller of a good both have relevant private information and cannot be forced to trade (i.e. participation is voluntary), there exists no incentive compatible ex post efficient trading mechanism. In settings where individuals' private information is multi-dimensional and their utility functions are interdependent (so that an individuals' private information directly affects the utility of other individuals), Jehiel and Moldovanu (2001) show that incentive compatible ex post efficient mechanisms exist only under a most restrictive condition; generically, such mechanisms do not exist.

A third important theme in Mechanism Design is the characterization of optimal

[4]See Mas-Colell et al. (1995) for a description of the so called *Expected Externality Mechanism.*

incentive compatible mechanisms - i.e. mechanisms which maximize some appropriately chosen notion of individual or collective welfare. For example, an optimal auction mechanism is one that maximizes the seller's expected welfare (i.e. revenue), where the welfare comparison of alternative procedures must be conducted at a time when the bidders' characteristics are still unknown. In symmetric settings, an optimal auction, conditional on the object being sold, is also ex post efficient.[5]

In this thesis, we follow up the above themes in three specific settings: In Chapter 2 we study the problem of designing optimal incentive contracts for information providers who, as part of their contractual obligations, come to possess private information. In the particular situation considered, a principal must choose between two alternatives with unknown payoffs. The information provider (or agent) initially has the same information about the alternatives as the principal, but, in contrast to the latter, can find them out by inspecting the alternatives at a fixed cost per alternative. Therefore, information should be acquired sequentially. The novel aspect of our model is that the alternatives are not ex ante identical, so that the order in which the alternatives are inspected matters.

Incentive problems in our model arise from two sources: First, the principal cannot observe whether the agent inspects any of the alternatives, and if so, in which order he inspects them. Second, the principal observes only the payoff associated with her chosen alternative, not that of the other. Contracts between principal and agent can be written on all observables, namely the agent's reports about the alternatives' payoffs, and the principal's actual payoff from her chosen alternative. We characterize optimal contracts that incentivize the expert to adopt the *first best* sequential inspection procedure - the one that the decision-maker would use if she could acquire information herself.

We show that the first best procedure is one of two types. Under the first type, all incentive problems can be resolved without granting the agent any rents (on average), provided that false reports about the principal's chosen alternative can be sanctioned. Sequential procedures of this type are associated with a fixed "threshold": if the first

[5]See Maskin (1992).

alternative has a payoff above the threshold, the principal will choose it straight away.

Under the second type of inspection procedure, the agent must be granted strictly positive rents even if false reports can be punished. The reason is that the principal's decision to choose one alternative over the other is endogenous, and hence the agent has an incentive to minimize inspection cost by inspecting only one alternative and claiming that the other one is worse. The rents that the agent can command may be so high that it is optimal for the principal to not to implement the first best inspection procedure. However, even if a different procedure can be implemented "costlessly", there is still an efficiency loss from agency despite full surplus extraction.

In Chapter 3 of this thesis we study a setting in which two agents have to choose one of three alternatives.[6] Their ordinal rankings are diametrically opposed to each other and are common knowledge among the agents. Ex ante efficiency requires that they implement the compromise, that is the alternative which they both rank second, if and only if the sum of their von Neumann Morgenstern utilities from this alternative exceeds the sum of their utilities from the extreme options. We suppose that the von Neumann Morgenstern utilities are privately observed types.

We ask whether there are incentive compatible mechanisms that elicit utilities and implement ex post efficient decisions when money transfers between the agents are ruled out.[7] In the absence of side payments, incentives for truthful revelation of private information must be given by exposing the agents to risk. In this setup, we derive an analogous principle to that underlying Revenue Equivalence Theorem for auctions: the expected utility of an agent is determined up to constant by the allocation rule that, for every possible type, indicates the probability with which the compromise is chosen.

Our main result is that there exist no incentive compatible ex post efficient mechanisms if the distribution of agents' types has continuum support. The proof of this impossibility bears some similarity to the proof of the impossibility of efficient bilateral

[6]Chapter 3 is joint work with Tilman Börgers.

[7]The problem we study is a simplified version of the voting problem because voting rules, if there are more than two candidates, must elicit information about voters' strength of preference for candidates in the absence of money transfers.

trade in Myerson and Satterthwaite (1983). We also show that in simple examples, in which there are only two or three types to which the prior distribution attaches positive values, there exist simple incentive compatible ex post efficient mechanisms.

Chapter 4 of this thesis is devoted to the design of procurement procedures in which not only the price of the good, but also its specification is determined as part of the process. We study the question whether procurement procedures that *simultaneously* determine specification and price of the good can result in an inefficient specification choice.

We consider a stylized model in which each of two suppliers can produce a good in either of two specifications, both of which are equally good for the buyer. Production costs are interdependent and unknown at the time of bidding. Each supplier has two-dimensional private information about the production cost.

We analyze a specific procurement mechanism that selects the winning supplier purely on the basis of price and irrespective of his chosen specification. Under this mechanism, there is a strictly positive chance that the specification with the highest production cost is selected.

While there exists an efficient and incentive compatible mechanism, it requires the buyer to make large lump-sum payments to the suppliers in order to give them incentives to truthfully reveal their cost information and to take part in the mechanism. Our setup is an example of the non-generic case in Jehiel and Moldovanu (2001) where incentive compatible ex post efficient mechanisms exists.

A comparison of the two mechanisms leads us to conjecture that in our setup with two-dimensional private information, an optimal mechanism (i.e. one that minimizes the buyer's expected expenditure for the good) will not generally be ex post efficient even if the buyer is obliged to purchase the good. This would contrast with the standard result on optimal auctions with scalar private information.

Chapter 2

Delegated Search

2.1 Introduction

In this chapter we study a contracting-model in which an expert can acquire and communicate information that helps a decision-maker in choosing the best available alternative. The real-world situations we have in mind include recruitment agencies who interview job candidates on behalf of a manager, media agencies who "identify" the best advertising campaign for their client's products, insurance brokers who find out which policy is best for a client, etc.

In our model, the expert can find out the payoffs of two risky alternatives. The decision-maker cannot observe whether the expert "inspects" the alternatives, and if so, which information this brings to light. All she observes is the payoff she receives from choosing an alternative, and by selecting one, she foregoes the possibility of learning the payoff of the other. Inspecting the alternatives is costly for the agent, and hence he has an incentive to acquire less information than requested.

The decision-maker's problem is to design a contract that induces the agent to investigate the alternatives. However, the decision-maker may not always want the agent to find out both payoffs. As there are no economies of scale from inspecting both alternatives at once, it is optimal to acquire information sequentially. The alternatives are not ex ante identical, and therefore the order in which they are inspected matters.

The purpose of this chapter is to characterize contracts that incentivize the expert to adopt the *first best* sequential inspection procedure - the one that the decision-maker would use if she could acquire information herself.

Depending on the model parameters, the first best procedure is one of two types. We show that one type can be "costlessly" delegated to a risk-neutral agent with limited liability, provided that false information, if detected, can be sanctioned. The other type, however, can only be implemented by granting the agent information rents, in which case it may be optimal to implement a different procedure.

The first procedure is associated with a payoff-threshold that each of the two alternatives can potentially meet or exceed. If the first alternative has a payoff above this threshold, the decision-maker chooses it straight away.[1] The optimal contract that implements this procedure pays the expert only if he succeeds in identifying an alternative with a payoff above the threshold. Otherwise the expert is not paid at all.[2]

Under the second type of inspection procedure the decision-maker chooses the first alternative only if the second one is worse. In this case, the agent must be reimbursed for the cost of two inspections. As the decision to choose the first over the second alternative is endogenous, it pays for the agent to investigate only the second alternative and pretend that the first one is worse.

Under the first type of procedure this behavior can be "costlessly" deterred because the cost of inspecting the first alternative are fully covered (in expected terms) by the promise of a generous payment if the alternative is chosen straight away. If the expert is subsequently asked to investigate the other alternative, he must be reimbursed only for the cost of the second inspection. It therefore does not pay for him to investigate only the second alternative in order to claim that he has seen both.

The present work makes a contribution to the literature on principal-agent problems

[1] If both alternatives fall short of the threshold the decision-maker has to make a compromise by selecting the "best of a bad lot".

[2] Contracts of this type bear some resemblance to the way in which advertising agencies are rewarded in reality: the agency is paid only if sales following the advertising-campaign exceed a preset target level.

with acquisition of soft information. The most general formulation of this problem is due to Demski and Sappington (1987). The basic structure is as follows: The principal must choose from amongst a range of available projects whose payoffs depend on the unknown state of the world. Thus, each project is a lottery over a given set of possible outcomes. The agent can, at a cost to himself, acquire a private signal about the state from a range of available signal technologies with varying precision. The higher the precision, the more costly the signal is to acquire. Demski and Sappington study the properties of optimal contracts for information acquisition and subsequent project selection in a series of examples.

Malcomson (2002) provides general results for a setup where the available projects are binary lotteries over the outcomes "success" and "failure". Projects differ with respect to the payoff that the principal obtains in the event of success, as well as their success probability. Malcomson shows that the principal may find it optimal to limit the information rents to a risk-neutral agent with limited liability by distorting the way in which information is used relative to the first-best level (i.e. relative to a situation in which the principal can acquire and use information herself). The present work is remotely related to Malcomson's in as far as the implementation of sequential information acquisition procedures requires the agent to appropriately use any information acquired in the first stage when deciding whether or not to acquire further information.

The paper by Martimort and Gromb (2003) is related to the present work via its focus on the implementation of (possibly) sequential information acquisition procedures. They, however, study a setting in which the agent can acquire two informative i.i.d. signals that convey information about the relative desirability of two available projects. Our model also features two signals, namely one per project. However, each signal is perfectly informative, and as the two projects are not identical, the signals are not i.i.d. Our model therefore features a richer set of incentive constraints because the order in which the signals are acquired matters.

A special feature of our model with perfectly informative signals is that the principal can compare the payoff from her chosen project with any information provided about it by the agent. We study the limiting case in which false reporting, if detected, can be punished. It is worth emphasizing, however, that this punishment assumption does *not* render the implementation of sequential information acquisition procedures trivial: Information rents arise from the agent's ability to acquire signals in whichever order he chooses.

A seminal paper that addresses issues surrounding delegation and information acquisition is that by Aghion and Tirole (1997). However, it has a different focus from the present one and the other papers mentioned above. Their paper could be described as dealing with information acquisition and delegation *within* organizations, as they consider an environment in which both principal *and* agent are affected by a decision to be taken. Principal and agent may differ in their view of what the "best" decision is. *Both* can acquire information about the available choices, and consequently, the principal can potentially check up on the agent. The main question is whether and when the principal should delegate decision-rights to the agent so as to provide incentives for information acquisition (which gives the agent the ability to take the decision that is best for him, but not necessarily for the principal). This, of course, has to be traded off against the potential loss of control from delegation for the principal. In our model, the principal cannot acquire information herself, and the agent is not directly affected by the choices available to the principal.

While dealing with an agent's incentives for information acquisition, the papers by Crémer and Khalil (1992) and Crémer, Khalil and Rochet (1998a) study a different problem. In their contracting environment, the agent can strategically acquire information that only directly affects himself, and which becomes freely available once he accepts the contract offered by the principal. Ideally, the principal would want the agent to remain uninformed. While studying two different variants of the same problem, both papers find that the possibility of becoming informed prior to accepting (or rejecting) the contract is a source of rents for the agent.[3] The paper by Crémer, Khalil

[3] In the paper by Crémer and Khalil (1992), the agent can only acquire information *after* a contract has been offered but *before* accepting it. In this case, information acquisition will optimally be deterred. The complementary paper by Crémer, Khalil and Rochet (1998*a*) studies the situation in which the

and Rochet (1998b) studies a third variant of this problem in which the agent will *not* know the relevant information upon accepting the contract unless he has acquired it between offer and acceptance. In this model, the principal trades off possible efficiency gains from dealing with an informed agent against the ability to inflict ex post losses on an uninformed agent.

The remainder of this chapter is structured as follows: In Section 2.2 we review the classical search model of Weitzman (1979) and introduce a simple extension that will constitute the basis of our delegated search model. Section 2.3 studies the hypothetical situation in which the principal can conduct interviews herself and characterizes the optimal inspection procedure. Section 2.4 contains the delegated search model as well as our main results. Section 2.5 concludes with a brief discussion. Longer proofs are relegated to Section 2.6 (Appendix to Chapter 2).

2.2 Search Model

In this and the following section, we focus on search problems involving a single decision-maker. These serve as benchmark against which we compare the results of delegated search once agency concerns are introduced:

A risk-neutral principal must choose between two alternatives indexed $i = 1, 2$. Each of them leads to a deterministic payoff that is unknown ex ante. The payoff associated with alternative i is x_i, which is the realization of a random variable $\widetilde{x}_i$ with domain $X_i \subset \mathbb{R}_+$, distribution G_i and mean μ_i. Let $\underline{x}_i$ and $\overline{x}_i$ respectively denote the minimal and maximal elements of X_i. Random variables $\widetilde{x}_1$ and $\widetilde{x}_2$ are statistically independent.

While the payoffs of the two alternatives are unknown ex ante, the principal can find them out before choosing by "inspecting" each one at a cost of c, where

$$0 < c \leq \min_{i \in \{1,2\}} \{E[\max\{\mu_i, \widetilde{x}_i\}] - \mu_i\}.$$

agent can acquire information in anticipation of (i.e. *before*) being offered a contract. In this case the agent may or may not acquire information, but obtains rents from the contract in either case.

When the principal inspects alternative 1 (2 resp.) she instantaneously observes the true payoff realization x_1 (resp. x_2). If she inspects both alternatives simultaneously, she instantaneously observes x_1 and x_2. As total inspection costs are additive in the number of inspections, it is never optimal to inspect the alternatives simultaneously. To make the issue of sequential information acquisition non-trivial, we exclude the following cases:

1. One alternative is always better than the other because it always yields a higher payoff than the other.

2. A single inspection suffices to reveal which of the two alternatives is best.

The principal's von Neumann Morgenstern utility function is $x_i - nc$, where x_i is the utility she obtains from choosing alternative i ($i = 1, 2$) and n is the number of inspections she makes. She faces the following search problem: Which alternative (if any) to inspect first, and, having inspected one, whether or not to find out the payoff of the other alternative. We consider two specifications of this problem:

In the first (referred to as "search with restricted choice"), the principal can only choose an alternative that has previously been inspected, as analyzed by Weitzman (1979).[4] For instance, if the principal has only inspected alternative 1, she cannot choose alternative 2. If she wants to choose between both alternatives, she must also inspect alternative 2. Prior to conducting any inspections, the principal can't choose either alternative. If she conducts no inspections, she is left with her outside option, assumed to yield a payoff of zero.

In the second specification of the problem (referred to as "search with unrestricted choice"), the principal is free to choose either alternative at any stage of the search process. Suppose she has inspected alternative 1. She may terminate search at this point by choosing alternative 2. Prior to conducting any inspections, she is free to choose either alternative 1 or 2 without incurring any search costs. In the following

[4] In Weitzman's (1979) model the decision-maker faces more than two alternatives with unknown payoffs.

section we use backward induction to characterize, for each of the two specifications of the search problem, the inspection procedure that maximizes the expected value of $x_i - nc$.

2.3 First Best Search Procedures

2.3.1 Search with restricted choice

What to do after the first inspection?

The principal's optimal search procedure must state which alternative to inspect first, and what to do once it has been inspected. Having inspected alternative i ($i = 1, 2$) the principal chooses the action that yields:[5]

$$\max\{x_i, E\left[\max\{x_i, \widetilde{x}_j\}\right] - c\},$$

where x_i is the payoff the principal gets from choosing alternative i, and $E\left[\max\{x_i, \widetilde{x}_j\}\right] - c$ is the expected payoff from inspecting alternative j at cost c before choosing the ex post optimal alternative. Figure 2.1 characterizes the principal's optimal contingent choice: Terminate search after the first inspection if x_i exceeds the "threshold" z_j. Otherwise, incur inspection cost c again in order to find out x_j before making a decision.

The "threshold" z_j is defined implicitly:

$$z_j = E\left[\max\{z_j, \widetilde{x}_j\}\right] - c. \tag{2.1}$$

This expression can be equivalently written as

$$c = \int_{z_j}^{\overline{x}_j} (x_j - z_j)\, dG_j(x_j). \tag{2.2}$$

This suggests the following interpretation: If the principal finds that alternative i's payoff is above z_j, the "option value" of being able to choose alternative j is lower than the "option price" c (which is the cost of observing x_j).[6]

[5]The cost of inspecting alternative i are omitted as they have already been sunk.

[6]Note that z_j need not be an element of X_i (for instance, $\widetilde{x}_i$ may be a discrete random variable).

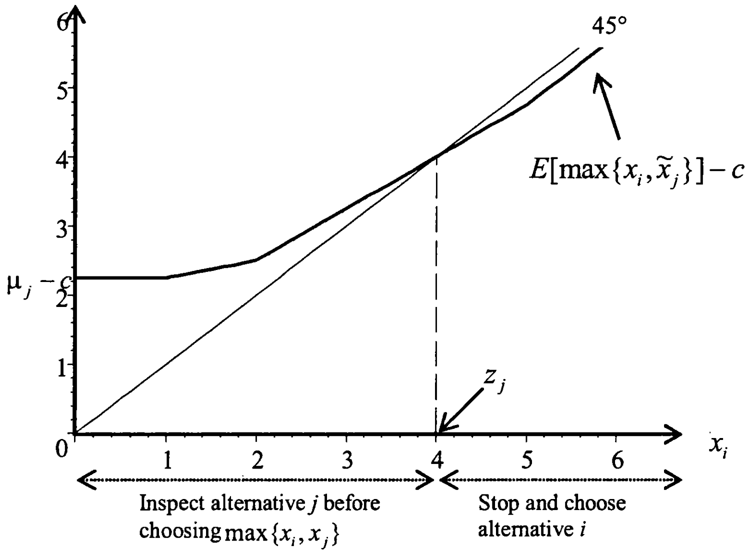

Figure 2.1: Optimal Contingent Choice (restricted)

Lemma 2.1 *For any alternative j ($j = 1, 2$) there exists a unique threshold z_j.*

Proof of Lemma 2.1: The result follows immediately from the properties of the function $E[\max\{\cdot, \widetilde{x}_j\}]$ and the fact that $c > 0$: If $z \leq \underline{x}_j$, then $E[\max\{z, \widetilde{x}_j\}] = \mu_j$. If $z \geq \overline{x}_j$, then $E[\max\{z, \widetilde{x}_j\}] = z$. Finally, $E[\max\{z, \widetilde{x}_j\}]$ is strictly increasing and convex for all $z \in [\underline{x}_j, \overline{x}_j]$.[7]

Q.E.D.

It should therefore be thought of as the *hypothetical* payoff that would make the principal indifferent between choosing alternative i and continuing the inspection process. Threshold z_j is the analogue of the Gittins Index in multi-armed bandit problems (see Roberts and Weitzman (1980)).

[7] As X_j has both a minimal and a maximal element, all elements of X_j lie in the interval $[\underline{x}_j, \overline{x}_j]$, which thereby constitutes the convex hull of X_j. As $E[\max\{z, \widetilde{x}_j\}] = \int_{\underline{x}_j}^{z} z dG_j(x_j) + \int_{z}^{\overline{x}_j} x_j dG_j(x_j)$, we have $\frac{d}{dz} E[\max\{z, \widetilde{x}_j\}] = G_j(z) > 0$.

Which alternative to inspect first?

Weitzman (1979) shows that the thresholds z_1 and z_2 not only determine the principal's optimal contingent choice once one alternative has been inspected, but also fully determine which alternative should be inspected first. Denote by π_i^r the expected of the sequential inspection procedure commencing with alternative i $(i = 1, 2)$ when the cost of the first inspection are disregarded.[8]

Proposition 2.1 (Weitzman (1979)) *$\pi_i^r \geq \pi_j^r$ if and only if $z_i \geq z_j$.*

The proof of Proposition 2.1 can be found in Weitzman's (1979) article. The following example provides an illustration:

Example *$X_1 = \{1, 2, 5\}$ and $X_2 = \{0, 3\}$. Let $\Pr(\widetilde{x}_1 = 1) = \Pr(\widetilde{x}_1 = 5) = 1/4$ and $\Pr(\widetilde{x}_2 = 0) = 1/2$ and $c = 1/4$. By equation (2.2) we have $z_1 = 4 > 5/2 = z_2$. It is easy to verify that $\pi_1^r = 45/16 > 44/16 = \pi_2^r$.*

2.3.2 Search with Unrestricted Choice

What to do after the first inspection?

If the principal can choose either alternative at any stage of the search process, the decision to stop searching after the first inspection does not automatically imply which alternative will be chosen. Having inspected alternative i $(i = 1, 2)$, the principal chooses the action that yields:

$$\max\{x_i, \mu_j, E[\max\{x_i, \widetilde{x}_j\}] - c\},$$

where x_i is the payoff the principal gets from choosing alternative i, μ_j is the expected payoff from choosing alternative j without inspection, and $E[\max\{x_i, \widetilde{x}_j\}] - c$ is the expected payoff from inspecting alternative j at cost c before choosing the ex post optimal alternative.

[8]The superscript r in π_i^r refers to "restricted choice".

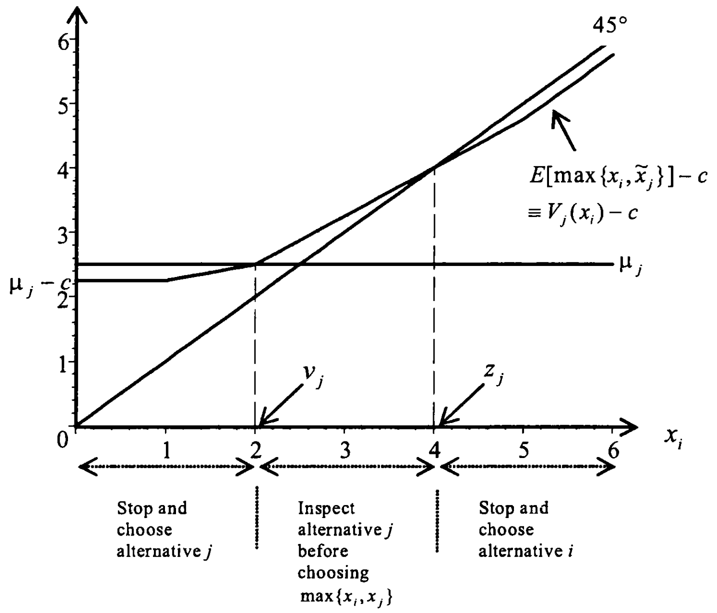

Figure 2.2: Optimal Contingent Choice (unrestricted)

Figure 2.2 characterizes the principal's optimal contingent choice: Terminate search after inspection of alternative i if x_i is "so good" that it exceeds the "upper threshold" z_j, or if x_i is "so bad" that it is below the "lower threshold" v_j. Otherwise, incur inspection cost c again in order to find out x_j before choosing the ex post optimal alternative.

The "upper threshold" z_j has been defined above in equation (2.1). The "lower threshold" v_j is the value that makes the principal indifferent between choosing j without inspection, and inspecting it before making a decision. Formally,

$$\mu_j = E\left[\max\left\{v_j, \tilde{x}_j\right\}\right] - c, \tag{2.3}$$

or

$$c = \int_{\underline{x}_j}^{v_j} (v_j - x_j)\, dG_j(x_j). \tag{2.4}$$

This suggests the following interpretation: If the principal finds a payoff realization x_i below v_j, then the value of "insurance" against a low payoff realization from alternative j is lower than the "insurance premium" c (which is the cost of observing x_j).

Lemma 2.2 *For any alternative j ($j = 1, 2$) there exists a unique threshold v_j.*

The proof follows immediately from that of Lemma 2.1 above. The assumption that $c \leq \min_{j=1,2}\{E\left[\max\left\{\mu_j, \widetilde{x}_j\right\}\right] - \mu_j\}$ ensures that $v_j \leq z_j$.

Which alternative to inspect first?

Denote by π_i^u the expected payoff from the sequential inspection procedure commencing with alternative i when the cost of the first inspection are disregarded:[9]

$$\pi_i^u = E[\max\{x_i, \mu_j, E\left[\max\left\{x_i, \widetilde{x}_j\right\}\right] - c\}], \text{ where } j = 1, 2, j \neq i.$$

It is easy to see from Figure 2.2 that $\pi_i^u \geq \max\{\mu_i, \mu_j\}$ for all i, so that it is always optimal to conduct at least one inspection. It is not possible to infer from the threshold values v_1, z_1, v_2 and z_2 alone which alternative should be inspected first. We can, however, derive a necessary condition for "order reversal", by which we mean that the optimal *inspection orders* under restricted and unrestricted choice differ.

Suppose that it is optimal to commence search with alternative i when choice is restricted: $\pi_j^r - \pi_i^r > 0$. Note that the expected payoff from the search procedure commencing with alternative j when choice is *un*restricted can be written as follows:

$$\pi_j^u = \pi_j^r + (\pi_j^u - \pi_j^r).$$

It is immediate from Figure 2.2 that $\pi_j^u - \pi_j^r \geq 0$. The payoff difference between the two search orders under unrestricted choice is

$$\pi_j^u - \pi_i^u = \pi_j^r - \pi_i^r + [(\pi_j^u - \pi_j^r) - (\pi_i^u - \pi_i^r)].$$

If $\pi_j^r - \pi_i^r > 0$ a necessary condition for order reversal is $(\pi_j^u - \pi_j^r) - (\pi_i^u - \pi_i^r) > 0$. We can show the following

Proposition 2.2 $(\pi_j^u - \pi_j^r) > (\pi_i^u - \pi_i^r)$ *if and only of $v_i > v_j$.*

The proof is relegated to the Appendix. If $\pi_i^r > \pi_j^r$ (which, by Proposition 2.1, holds if and only if $z_i > z_j$), then $v_i > v_j$ is necessary for the optimal inspection procedure

[9]The superscript u in π_i^u stands for "unrestricted choice".

under unrestricted choice to start with alternative j. If, instead, $v_i > v_j$ and $z_j \geq z_i$, it follows from Propositions 2.1 and 2.2 that

$$\pi_j^u - \pi_i^u > \pi_j^r - \pi_i^r \geq 0.$$

In this case the optimal inspection order is the same under restricted and unrestricted choice. The following example illustrates a situation where order reversal occurs:

Example cont'd: $X_1 = \{1, 2, 5\}$ *with* $\Pr(\widetilde{x}_1 = 1) = \Pr(\widetilde{x}_1 = 5) = 1/4$ *and* $X_2 = \{0, 3\}$ *with* $\Pr(\widetilde{x}_2 = 0) = 1/2$. *Also,* $c = 1/4$. *From equations (2.2) and (2.4) we obtain* $v_1 = 2$, $z_1 = 4$, $v_2 = 1/2$ *and* $z_2 = 5/2$. *With restricted choice it is optimal to first inspect alternative 1, while with unrestricted choice it is optimal to first inspect alternative 2* ($\pi_2^u - \pi_1^u = 1/16$)*. The intuition for order reversal is as follows: As* $v_2 < \underline{x}_1 = 1$ *the choice restriction imposes no "loss" when search commences with alternative 1 (i.e.* $\pi_1^u = \pi_1^r = 45/16$*). If search starts with alternative 2 and* $x_2 = 0$ *the principal knows that alternative 1 has a higher payoff. By choosing alternative 1 without inspection the principal "saves" inspection cost of* $1/4$ *that she would have to incur if choice was restricted. Thus,* $\pi_2^u - \pi_2^r = 1/8$.

2.4 Delegated Search

2.4.1 Setup

In this section we study the contractual relationship between the principal and an agent who can acquire information on her behalf. The principal is faced with the choice problem described in Section 2.2, but now cannot conduct inspections herself. By choosing an alternative, the principal instantaneously learns its payoff, but thereby foregoes the possibility of learning the payoff of the other alternative. The principal can learn the alternatives' payoffs prior to choosing by employing the agent. The agent can inspect any alternative by incurring cost c per alternative. If the agent inspects alternative 1 (resp. 2) he instantaneously observes the true realization x_1 (resp. x_2). If he inspects both alternatives simultaneously, he instantaneously observes x_1 and

x_2. Both principal and agent have symmetric information ex ante, that is, the sets X_1 and X_2, as well as the distribution functions G_1 and G_2 are common knowledge among them. The agent's effort cost c is also common knowledge. The principal's von Neumann-Morgenstern utility function is $x_i - w$, where x_i is the utility she obtains by choosing alternative i ($i = 1, 2$), and w is the payment she makes to the agent. The agent's von Neumann-Morgenstern utility function is $w - nc$, where n is the number of inspections he carries out. Both principal and agent are risk-neutral.

As a benchmark, consider the hypothetical scenario in which the principal completely controls the agent. Whatever instructions the principal gives to the agent, the agent will carry out. Assuming perfectly transferable utility, this situation is equivalent to the one considered in Section 2.3 where the principal carries out the inspections herself so as to maximize the expected value of $x_i - nc$. The inspection procedure that is optimal in this hypothetical situation is called *first best.*

The agent is free to choose his actions and the messages that he reports back to the principal. As a consequence, there is both a *hidden action problem* (will the agent carry out the appropriate number and order of inspections?) and a *hidden information problem* (will the agent report what he finds correctly?). The principal must therefore offer the agent a contract that addresses these incentive problems.

We model the contract negotiation between principal and agent as an ultimatum bargaining game: The principal proposes a contract, and the agent can either accept or reject it. If he rejects it, he obtains a utility of zero. Formally, a contract is defined as follows:

Definition 2.1 (Contract) *A contract is a triplet (i, f, w), where i is the alternative that the principal wants the agent to inspect first; the mapping f denotes the choice rule that states which alternative the principal chooses given the agents report(s): the agent either submits a singleton report about the payoff of alternative i, or he submits a report consisting of two payoffs - one about alternative i and one about alternative j; the mapping w determines the agent's wage as a function of his report and any payoff level that the principal observes upon choosing an alternative.*

We assume that a contract of this type is enforceable. In addition, we make the following assumptions:

Assumption 2.1 (No Transfers from Agent to Principal) *The agent has no initial endowment of money and cannot borrow.*

Assumption 2.1 implies that w must be non-negative. Its purpose is to ensure that the contracting problem is not trivial. In particular, it prevents the principal from solving all incentive problems by letting the agent choose whichever alternative he wants and transferring its payoff to him in exchange for a fixed payment.[10] It is worth emphasizing that this assumption does *not* require the agent's ex post utility $w - nc$ to be non-negative in all possible outcomes that may arise under the terms of the contract. The agent's (net) utility $w - nc$ may be negative in some outcomes. However, in order to ensure acceptance of the contract the agent's *expected* utility from it must be at least zero. To avoid inconsistencies with Assumption 2.1, we assume henceforth that the agent's inspection cost c are non-monetary.

Assumption 2.2 (Punishment) *If the payoff of the principal's chosen alternative differs from the agent's report about it, the principal imposes an arbitrarily large non-monetary punishment on the agent.*

One can interpret the punishment as irreversible damage to the agent's professional reputation. Assumption 2.2 prevents the agent from submitting reports for which the principal chooses an alternative whose payoff differs from his report with positive probability.

2.4.2 Costless Implementation

In this section we characterize inspection procedures that can be delegated to the agent without granting him information rents. We derive a simple contract that extracts, on

[10] By charging a negative wage equal to the principal's expected payoff from the first best inspection procedure the agent is induced to adopt this procedure, while all his (expected) rents are extracted.

average, all surplus from the agent. Suppose the principal wishes to implement a sequential inspection procedure that starts with alternative i ($i = 1, 2$) and which satisfies the following

Condition 2.1 $z_j < \overline{x}_i$, *where* $j = 1, 2$, $j \neq i$.

This condition says that there exist payoff realizations x_i that are "so good" that the principal chooses alternative i without further information about alternative j. Figure 2.3 shows such a procedure under unrestricted choice.[11] The interval $[\underline{x}_i, \overline{x}_i]$ depicted in the figure represents the set of possible payoff realizations of alternative i. For values x_i in $[\underline{x}_i, v_j]$ the principal chooses alternative j without inspection. For payoff realizations x_i in $(v_j, z_j]$ there is valuable information to be gained from an inspection of alternative j. Finally, for payoff realizations x_i in $(z_j, \overline{x}_i]$ the principal chooses alternative i immediately.

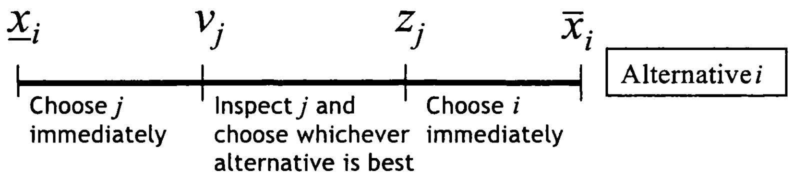

Figure 2.3: Search Procedure that satisfies Condition 2.1

In the following, we propose a simple contract (i, f^*, w^*) that allows the principal to implement inspection procedures that satisfy Condition 2.1. It asks the agent to first inspect alternative i. The principal's decision rule f^* replicates the choices associated with the inspection procedure shown in Figure 2.3:[12]

$$f^* = \begin{cases} i & \text{if } \widehat{x}_i > z_j \vee (v_j \leq \widehat{x}_i \leq z_j \wedge \widehat{x}_i > \widehat{x}_j) \\ j & \text{if } \widehat{x}_i \leq v_j \vee (v_j \leq \widehat{x}_i \leq z_j \wedge \widehat{x}_i \leq \widehat{x}_j) \end{cases}, \tag{2.5}$$

[11]For what follows it is irrelevant whether choice is restricted or unrestricted.

[12]Given the inspection procedure depicted in Figure 2.3, the domain of the mapping f^* is the set

$$([\underline{x}_i, v_j] \cup [z_j, \overline{x}_i]) \cup (\{x_i \,| v_j < x_i < z_j\} \times X_j).$$

where $\widehat{x}_i$ is the agent's report about the payoff of alternative i, and $\widehat{x}_j$ is his report about the payoff of alternative j. The wage-schedule is given by:

$$w^* = \begin{cases} \frac{c}{1-G_i(z_j)} & \text{if } \widehat{x}_i = x_i > z_j \\ \frac{c}{1-G_j(z_j)} & \text{if } v_j \leq \widehat{x}_i \leq z_j \text{ and } \widehat{x}_j = x_j > z_j \\ 0 & \text{otherwise} \end{cases} . \tag{2.6}$$

The interpretation is as follows: The agent receives a positive payment only in case of "success", i.e. if the payoff from the principal's chosen alternative exceeds the upper threshold z_j. If alternative i is chosen and its payoff is above this threshold, then the agent receives a fixed payment of $c/\left(1 - G_i(z_j)\right)$. If alternative j is chosen and its payoff is above the threshold, then the agent receives a fixed payment of $c/\left(1 - G_j(z_j)\right)$. In all other cases the agent receives no payment. The lump-sum payments in w^* in (2.6) are chosen so as to make the agent's expected payment for any inspection equal to the inspection cost. Thus, at any stage of the search process, the agent is indifferent between conducting and not conducting the inspection stipulated by the contract.

Proposition 2.3 (Costless Implementation) *Under contract (i, f^*, w^*), where f^* is given by (2.5) and w^* is given by (2.6), the agent's optimal strategy is to inspect alternative i first. If $x_i \leq v_j$ or $x_i > z_j$ the agent terminates search and submits the report $\widehat{x}_i = x_i$. If $v_j < x_i \leq z_j$ the agent inspects alternative j and submits the report $(\widehat{x}_i, \widehat{x}_j) = (x_i, x_j)$.*

The agent's expected payoff from contract (i, f^*, w^*) is zero, and the principal's expected payoff is the same as in the first best. This result rests on the assumption that whenever the agent is indifferent between several "actions" he will adopt the one preferred by the principal.[13]

Proof of Proposition 2.3: We show that under (i, f^*, w^*) the agent cannot obtain a positive expected payoff by deviating from the principal's desired inspection procedure:

[13]It is easy to modify the wage-schedule in (2.6) so as to give the agent strict incentives to follow the principal's desired procedure: If alternative i is chosen and has payoff above z_j, pay the fixed wage $\frac{c}{1-G_i(z_j)} + \varepsilon$, and if alternative j is chosen and has payoff above z_j, pay the fixed wage $\frac{c}{1-G_j(z_j)} + \delta$, where $\varepsilon > \delta(1 + \frac{G_i(v_j)}{1-G_i(z_j)}) > 0$.

(i) **No inspections**: If the agent accepts the contract but inspects neither alternative, he must submit a report $\widehat{x}_i \leq v_j$ in order to avoid the risk of punishment. In this case the principal chooses alternative j and therefore cannot verify the agent's report about alternative i. The agent's wage in this case is zero.

(ii) **Inspecting alternative j first**: If the agent finds that $x_j > z_j$, he can secure a wage of $c/\left(1 - G_j(z_j)\right)$ by submitting any report $(\widehat{x}_i, x_j)$ where $v_j < \widehat{x}_i \leq z_j$. This allows him to pretend that both alternatives have been inspected. If $x_j \leq z_j$ the agent can either terminate the inspection process or go on to inspect alternative i "retrospectively". In the former case the highest wage he can secure is 0 (e.g. by submitting any report $\widehat{x}_i \leq v_j$). It is easy to verify that the expected payoff from "retrospectively" inspecting alternative i is also zero. Thus, the agent's ex ante wage from any inspection procedure that commences with alternative j is zero.

(iii) **Inspecting alternative i only**: Suppose the agent inspects alternative i. If $x_i > z_j$ he can secure a wage of $c/\left(1 - G_i(z_j)\right)$ by truthfully reporting x_i. If $x_i \leq z_j$ the highest wage the agent can get is zero (e.g. by submitting any report $(x_i, \widehat{x}_j)$ for which $\widehat{x}_j < x_i$). Thus, the agent's ex ante wage from inspecting alternative i only is zero.

(iv) **Inspecting i and j simultaneously**: In this case the agent's ex ante wage is $-c(1 - G_i(z_j) + G_i(v_j)) < 0$.

Q.E.D.

An implication of Proposition 2.3 is that *if* the principal's first best inspection procedure satisfies Condition 2.1, then contract (i, f^*, w^*) is optimal. This follows immediately from the fact that the principal's expected payoff is the same as in the first best. It is easy to show the following

Proposition 2.4 *The first best inspection procedure under restricted choice satisfies Condition 2.1.*

Proof of Proposition 2.4: Suppose that it is optimal to first inspect alternative i. Thus, $z_i > z_j$. We have to show that $z_i > z_j$ implies $\overline{x}_i > z_j$. To see that this is true recall that $\overline{x}_i > z_i$ since $c > 0$. This immediately implies that $\overline{x}_i > z_i > z_j$.

Q.E.D.

Thus, if choice is restricted, the first best can always be attained by contract (i, f^*, w^*). It is also worth noting that any inspection procedure that satisfies Condition 2.1 can also be costlessly implemented if the principal cannot write "contingent" contracts. When restricted to contracts that cover only the inspection of a single alternative, the principal should first offer the agent a contract (i, w_1) to inspect alternative i, where the wage is determined by

$$w_1 = \begin{cases} \frac{c}{1-G_i(z_j)} & \text{if } \widehat{x}_i > z_j \\ 0 & \text{otherwise} \end{cases} .$$

Based on the report $\widehat{x}_i$ and the belief that $\widehat{x}_i = x_i$, the principal chooses the action which is ex post optimal: If $\widehat{x}_i > z_j$, the principal chooses alternative i. If $\widehat{x}_i \leq v_j$, she chooses alternative j. If $v_j < \widehat{x}_i \leq z_j$, she offers the agent the contract (j, w_2) to inspect alternative j. The function w_2 is given by:

$$w_2 = \begin{cases} \frac{c}{1-G_j(z_j)} & \text{if } \widehat{x}_j > z_j \\ 0 & \text{otherwise} \end{cases} .$$

Once the agent has inspected both alternatives, the principal chooses the one which yields $\max\{\widehat{x}_i, \widehat{x}_j\}$, based on her belief that both reports are correct.

2.4.3 Costly Implementation

In this section, we show that if choice is unrestricted, so that the principal can choose any alternative irrespective of whether or not the agent has previously inspected it, then it is not possible to costlessly delegate first best inspection procedures that violate Condition 2.1. Suppose the first best procedure commences with alternative i and violates Condition 2.1. This implies that the first best search order under restricted and unrestricted choice differ.[14] We know from Proposition 2.2 that "order reversal"

[14]This follows immediately from the violation of Condition 2.1: we have $z_j \geq \overline{x}_i$ and hence $z_j > z_i$.

implies that $v_j > v_i$. Therefore any inspection procedure that is first best under unrestricted choice and violates Condition 2.1 is of the form depicted in Figure 2.4. It depicts the possible payoff realizations of alternatives i and j as overlapping subsets of the real line. The key feature of such search procedures is that alternative i (the one to be inspected first) is chosen if and only if alternative j has a lower payoff. Thus, the decision to choose i relies on a direct comparison of alternatives i and j.

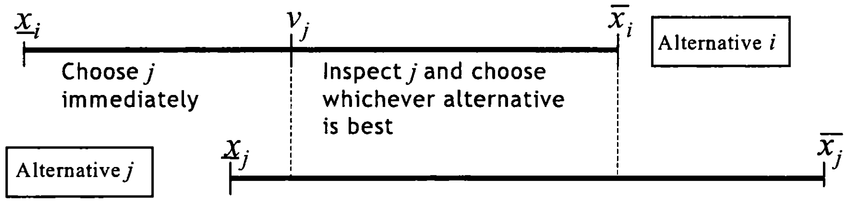

Figure 2.4: Inspection Procedure that violates Condition 2.1

A generic incentive contract (i, f, w) that implements an inspection procedure of the type depicted in Figure 2.4 has the following form:[15]

$$f = \begin{cases} i & \text{if } \widehat{x}_i > \max\{v_j, \widehat{x}_j\} \\ j & \text{if } \widehat{x}_i \leq v_j \vee (v_j < \widehat{x}_j < \widehat{x}_i) \end{cases} . \tag{2.7}$$

The agent's payment is determined by a wage-schedule of the following form:

$$w = \begin{cases} w_0(x_j) & \text{if } \widehat{x}_i \leq v_j \\ w_1(x_i) & \text{if } \widehat{x}_i > v_j \text{ and } \widehat{x}_i = x_i > \widehat{x}_j \\ w_2(x_j) & \text{if } \widehat{x}_i > v_j \text{ and } \widehat{x}_j = x_j > \widehat{x}_i \end{cases} \tag{2.8}$$

We can now establish the following result

Proposition 2.5 (Impossibility of costless implementation) *A first best inspection procedure that violates Condition 2.1 can only be implemented by granting the agent strictly positive rents.*

[15]Given the inspection procedure depicted in Figure 2.4, the domain of the mapping f is the set

$[\underline{x}_i, v_j] \cup ([v_j, \overline{x}_i] \times X_j)$.

Proof of Proposition 2.5: By contradiction. Suppose there exists a wage schedule w of the form shown in (2.8) that is accepted by the agent and costlessly implements the inspection procedure depicted in Figure 2.4. This implies that no other inspection procedure yields the agent a positive expected payoff:

(i) **No inspections**: If the agent accepts the contract but inspects neither alternative, he must submit a report $\widehat{x}_i \leq v_j$ in order to avoid the risk of punishment. In this case the principal chooses alternative j and therefore cannot verify the agent's report about alternative i. This deviation cannot yield a positive wage. Thus it must hold that $w_0(x_j) = 0$ for all x_j.

(ii) **Inspecting alternative i only**: Suppose the agent has inspected alternative i and has found the payoff realization x_i, where $x_i > v_j$. If the agent deviates from the principal's desired inspection procedure and terminates the inspection process, the highest wage he can secure himself is $w_1(x_i)$.

As a benchmark, consider the agent's expected payoff from conducting the second inspection as required by the contract: the agent is paid $w_1(x_i)$ if $x_j \leq x_i$ and $w_2(x_j)$ if $x_j > x_i$. Thus, the expected payoff from the second inspection is

$$w_1(x_i)G_j(x_i) + \int_{x_i}^{\overline{x}_j} w_2(x_j)dG_j(x_j) - c.$$

As the contract is assumed to be incentive compatible, it cannot pay for the agent to terminate search after the first inspection if it reveals that $x_i > v_j$. Thus,

$$w_1(x_i) \leq w_1(x_i)G_j(x_i) + \int_{x_i}^{\overline{x}_j} w_2(x_j)dG_j(x_j) - c,$$

or equivalently

$$\int_{x_i}^{\overline{x}_j} w_2(x_j)dG_j(x_j) - c \geq (1 - G_j(x_i))w_1(x_i). \tag{2.9}$$

(iii) **Inspecting alternative j first**: If the agent finds that $x_j > v_j$, he can truthfully reveal x_j while claiming that alternative i has a lower payoff (i.e. submit any report $\widehat{x}_i < x_j$). In this case, the principal chooses alternative j and the agent's wage is $w_2(x_j)$.[16] If, instead, the agent finds a payoff $x_j \leq v_j$ he can either terminate search, or "retrospectively" inspect i. In the former case he must pretend that he has inspected alternative i but that its payoff is low (i.e. $\widehat{x}_i \leq v_j$). From (i) we know that such reports yield a wage of zero. In the latter case, the agent's expected wage from the second inspection is

$$\int_{v_j}^{\overline{x}_i} w_1(x_i)dG_i(x_i) - c.$$

The ex ante payoff from starting with alternative j therefore is

$$G_j(v_j)\max\{0, \int_{v_j}^{\overline{x}_i} w_1(x_i)dG_i(x_i) - c\} + \int_{v_j}^{\overline{x}_j} w_2(x_j)dG_j(x_j) - c.$$

As the contract is incentive compatible, this inspection procedure cannot yield a positive payoff:

$$G_j(v_j)\max\{0, \int_{v_j}^{\overline{x}_i} w_1(x_i)dG_i(x_i) - c\} + \int_{v_j}^{\overline{x}_j} w_2(x_j)dG_j(x_j) - c \leq 0. \tag{2.10}$$

Note that the left-hand side of (2.10) is at least as large as the left-hand side of (2.9), which, in turn, is non-negative. Thus, both the left-hand side of (2.9) and that of (2.10) must be equal to zero, which (by (2.9)) implies that $w_1(x_i) = 0$ for all $x_i > v_j$. Hence,

$$\max\{0, \int_{v_j}^{\overline{x}_i} w_1(x_i)dG_i(x_i) - c\} = 0.$$

We therefore obtain

$$\int_{x_i}^{\overline{x}_j} w_2(x_j)dG_j(x_j) = c = \int_{v_j}^{\overline{x}_j} w_2(x_j)dG_j(x_j).$$

The latter equality is satisfied if and only if $w_2(x_j) = 0$ for all $x_j \in [v_j, \overline{x}_i]$. The agent is therefore not paid at all if alternative i is chosen. Also, the agent's wage if alternative j is chosen just covers (in expected terms) the inspection cost. This implies that the expected wage from inspecting alternative j, conditional on having found that alternative i's payoff is above v_j, is the same as the unconditional expected wage (which the agent can get by inspecting alternative j only).

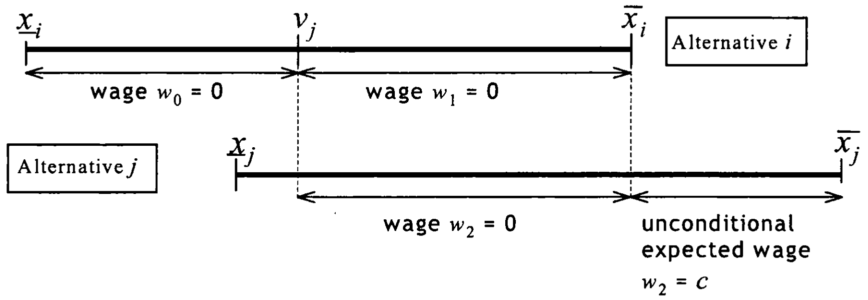

Figure 2.5: Contract violates Participation Constraint

It is easy to see that the above incentive compatible wage-schedule violates the agent's participation constraint (Figure 2.5 provides an illustration): The agent is reimbursed (in expected terms) only for the second inspection. Thus, the expected payoff from accepting the contract is $-c$ (which is the cost of conducting the first inspection). This constitutes a contradiction to the premise that the agent is willing to accept the contract.

Q.E.D.

The following proposition characterizes a contract that implements the first best inspection procedure depicted Figure 2.4 at the lowest cost to the principal.

Proposition 2.6 *Contract (i, f, w), where f is given by (2.7) and*

$$w = \begin{cases} 0 & \text{if } x_i \leq v_j \text{ and } x_j \leq \overline{x}_i \\ \underline{w} & \text{if } x_i > v_j \text{ and } x_j \leq \overline{x}_i \\ \overline{w} & \text{if } x_i \leq v_j \text{ and } x_j > \overline{x}_i \\ \underline{w} + \frac{c}{1-G_j(\overline{x}_i)} & \text{if } x_i > v_j \text{ and } x_j > \overline{x}_i \end{cases},$$

where

$$\underline{w} = \frac{(G_i(v_j) + (1 - G_i(v_j))(1 - G_j(v_j)))\, c}{G_j(v_j)G_i(v_j)(1 - G_i(v_j))}$$

and

$$\overline{w} = \frac{(1 - G_j(v_j))c}{G_i(v_j)(1 - G_j(\overline{x}_i))} + \frac{(1 - G_j(v_j))}{(1 - G_j(\overline{x}_i))}\underline{w}$$

[16] As the report $\widehat{x}_i$ on alternative i cannot be verified, the wage w_2 must not depend on it.

minimizes the expected wage that the principal has to pay in order to implement inspection procedures that violate Condition 2.1.

The proof is relegated to the Section 2.6 (Appendix). The contract in Proposition 2.6 yields the agent a strictly positive expected utility. To see this, suppose that he deviates from the desired inspection procedure and conducts no inspections at all. In order to avoid punishment, he must submit some report $\widehat{x}_i \leq v_j$. In this case the principal chooses alternative j straight away and the agent receives a strictly positive wage if the principal's payoff from alternative j is above $\overline{x}_i$. Thus, the agent's ex ante payoff from conducting no inspections is $(1 - G_j(\overline{x}_i))\overline{w}$.

An obvious question in cases where the first best inspection procedure does not satisfy Condition 2.1 is whether it is optimal for the principal to implement the first best procedure, or whether it is better to costlessly implement the "reverse" search procedure. The following example shows that the latter may indeed be optimal for the principal.

Example cont'd: *As before,* $X_1 = \{1, 2, 5\}$ *with* $\Pr(\widetilde{x}_1 = 1) = \Pr(\widetilde{x}_1 = 5) = 1/4$ *and* $X_2 = \{0, 3\}$ *with* $\Pr(\widetilde{x}_2 = 0) = 1/2$. *Also,* $c = 1/4$. *Choice is unrestricted and hence the first best inspection procedure commences search with alternative* 2. *The expected payoff from this procedure is* 2.625. *If it is delegated to the agent using the wage schedule in (2.29), the agent faces a wage of* 0 *for reports/outcomes* $(x_2, x_1) = (0, 1)$ *and* $(0, 2)$. *For reports/outcomes* $(3, 1)$ *and* $(3, 2)$ *he gets a wage* $\underline{w} = 5/6$, *for* $(0, 5)$ *he gets* $\overline{w} = 8/6$, *and for* $(3, 5)$ *he gets* $11/6$. *This leaves him an ex ante utility of* $1/3$. *Consequently, the principal's payoff when implementing the first best inspection procedure is (just under)* 2.292. *In contrast, the inspection procedure commencing with alternative 1 does satisfy Condition 1 (the principal chooses alternative 1 after the first inspection if its payoff is reported to be* 5*). It can therefore be delegated to the agent without granting him rents, yielding the principal an expected payoff of* $2.5625 > 2.292$.

2.5 Discussion

The present work shows that whether a sequential search procedure can be costlessly delegated to an information provider depends on the form of the procedure. If there is a chance that, after every single report made by the agent, the principal will pick the corresponding alternative, then the agent can be costlessly induced to follow the desired procedure. The reason is that every single inspection potentially generates a "verifiable" outcome to which the principal can tie her payment. Therefore, the expected payoff from any single inspection can always be lowered to zero. For inspection procedures where the first alternative is chosen if and only if the second one is worse, the principal's decision to choose the first one is endogenous. In this case, the agent must be granted information rents in order to follow the principal's desired procedure. We have shown by means of an example that the agent's rents can be sufficiently large, so that it is better for the principal to use a different inspection procedure. Even if the alternative procedure can be implemented costlessly, there is an efficiency loss because the implemented search procedure is not the first best one.

Our results suggest that in real life recruitment we should observe widespread use of recruitment agents when it comes to filling job-openings for which it is easy to fix a "performance standard" or objective selection criterion. Casual observation suggests that is indeed so, with companies making frequent use of recruitment agencies to hire support and administrative staff. Note also that recruitment agents are typically paid for their services only if they successfully place a candidate, in which case the agency receives a fixed fee that is a proportion of the position's salary. In this sense, real-life contracts for recruitment agents (at least in high-volume recruitment) carry some hallmarks of the contract derived here. The fact that recruitment agencies often make recourse to candidate-data-bases instead of newly searching for potential candidates could be seen as a response over time to the type of contract proposed here, as it exposes the agent to losses even if he has interviewed candidates, but fails to make the placement.

2.6 Appendix

Proof of Proposition 2.2: Suppose that $v_i > v_j$. We need to show that this implies $(\pi_j^u - \pi_j^r) - (\pi_i^u - \pi_i^r) > 0$. We first consider the inspection procedure commencing with alternative i and compute its expected payoff in the case of unrestricted and restricted choice (π_i^u and π_i^r resp.). The two cases differ only for values $x_i \le v_j$. We can therefore write (where the term "Rest" is used to describe all elements that are present both in π_i^u and π_i^r):

$$\pi_i^u = G_i(v_j)\mu_j + \text{Rest} \tag{2.11}$$

and

$$\begin{aligned}\pi_i^r \;=\;& G_i(v_j)(G_j(v_j)E[\max\{x_i, x_j\}\,|x_i \le v_j, x_j \le v_j] \\ &+(1 - G_j(v_j))E[x_j\,|x_j > v_j] - c) + \text{Rest}.\end{aligned} \tag{2.12}$$

Subtracting (2.12) from (2.11) yields after some manipulation

$$\pi_i^u - \pi_i^r = G_i(v_j)[c - G_j(v_j)(E[\max\{x_i, x_j\}\,|x_i \le v_j, x_j \le v_j] - E[x_j\,|x_j \le v_j])]. \tag{2.13}$$

The expression in (2.13) can be simplified further by using the fact that $c = \int_{\underline{x}_j}^{v_j} (v_j - x_j)\, dG_j(x_j)$, which implicitly defines v_j. It can obviously be rewritten as $c = G_j(v_j)(v_j - E[x_j\,|x_j \le v_j])$. Substituting this expression into (2.13) finally yields

$$\pi_i^u - \pi_i^r = G_i(v_j)G_j(v_j)[v_j - E[\max\{x_i, x_j\}\,|x_i \le v_j, x_j \le v_j]]. \tag{2.14}$$

This expression is non-negative.

We now consider the inspection procedure commencing with alternative j and compute its expected payoff under restricted and unrestricted choice. The two cases differ only for values $x_j \le v_i$. Taking into account that $v_j < v_i$, we can therefore write (where the term "Rest" is again used to describe all elements that are present both in π_j^u and π_j^r):

$$\pi_j^u = G_j(v_j)\mu_i + (G_j(v_i) - G_j(v_j))\,\mu_i + \text{Rest}, \tag{2.15}$$

and

$$\begin{aligned}
\pi_j^r &= G_j(v_j)\{G_i(v_j)E[\max\{x_i,x_j\}\,|x_i \le v_j, x_j \le v_j] \\
&\quad +(G_i(v_i)-G_i(v_j))E[x_i\,|v_j \le x_i \le v_i] \\
&\quad +(1-G_i(v_i))E[x_i\,|x_i > v_i] - c\} \\
&\quad +(G_j(v_i)-G_j(v_j))\{G_i(v_j)E[x_j\,|v_j \le x_j \le v_i] \\
&\quad +(G_i(v_i)-G_i(v_j))E[\max\{x_i,x_j\}\,|v_j \le x_i \le v_i, v_j \le x_j \le v_i] \\
&\quad +(1-G_i(v_i))E[x_i\,|x_i > v_i] - c\} + \text{Rest.} \qquad (2.16)
\end{aligned}$$

Since

$$\begin{aligned}
\mu_i &= G_i(v_j)E[x_i\,|x_i \le v_j] + (G_i(v_i)-G_i(v_j))E[x_i\,|v_j \le x_i \le v_i] \\
&\quad +(1-G_i(v_i))E[x_i\,|x_i > v_i],
\end{aligned}$$

we can express the difference between (2.15) and (2.16) as follows:

$$\begin{aligned}
\pi_j^u - \pi_j^r &= G_j(v_j)\{G_i(v_j)E[x_i\,|x_i \le v_j] - G_i(v_j)E[\max\{x_i,x_j\}\,|x_i \le v_j, x_j \le v_j] + c\} \\
&\quad +(G_j(v_i)-G_j(v_j))\{G_i(v_j)E[x_i\,|x_i \le v_j] - G_i(v_j)E[x_j\,|v_j \le x_j \le v_i]\} \\
&\quad +(G_i(v_i)-G_i(v_j))E[x_i\,|v_j \le x_i \le v_i] \\
&\quad -(G_i(v_i)-G_i(v_j))E[\max\{x_i,x_j\}\,|v_j \le x_i < v_i, v_j \le x_j \le v_i] + c\}. \quad (2.17)
\end{aligned}$$

Now note that the equation which implicitly defines v_i, $c = \int_{\underline{x}_j}^{v_i}(v_i - x_i)dG_i(x_i)$, can be written equivalently as

$$c = G_i(v_j)(v_i - E[x_i\,|x_i \le v_j]) + (G_i(v_i) - G_i(v_j))(v_i - E[x_i\,|v_j \le x_i \le v_i]).$$

Using this expression to eliminate c, we obtain

$$\begin{aligned}
\pi_j^u - \pi_j^r &= G_i(v_i)G_j(v_j)[v_i - E[x_i\,|v_j \le x_i \le v_i]] \qquad (2.18) \\
&\quad +G_j(v_j)G_i(v_j)[E[x_i\,|v_j \le x_i \le v_i] - E[\max\{x_i,x_j\}\,|x_i \le v_j, x_j \le v_j]] \\
&\quad +(G_j(v_i)-G_j(v_j))G_i(v_j)[v_i - E[x_j\,|v_j \le x_j \le v_i]] \\
&\quad +(G_j(v_i)-G_j(v_j))(G_i(v_i)-G_i(v_j))[v_i - E[\max\{x_i,x_j\}\,|v_j \le x_i \le v_i, v_j \le x_j \le v_i].
\end{aligned}$$

This expression is non-negative. To complete the proof, subtract (2.14). This yields the following expression for $(\pi_j^u - \pi_j^r) - (\pi_i^u - \pi_i^r)$:

$$\begin{aligned}
&G_i(v_i)G_j(v_j)[v_i - E[x_i\,|v_j \le x_i \le v_i]] + G_i(v_j)G_j(v_j)[E[x_i\,|v_j \le x_i \le v_i] - v_j] \\
&+(G_j(v_i) - G_j(v_j))G_i(v_j)[v_i - E[x_j\,|v_j \le x_j \le v_i]] \qquad (2.19)\\
&+(G_j(v_i) - G_j(v_j))(G_i(v_i) - G_i(v_j))[v_i - E[\max\{x_i, x_j\}\,|v_j \le x_i \le v_i, v_j \le x_j \le v_i]].
\end{aligned}$$

As $v_i > v_j$, it is easy to see that $(\pi_j^u - \pi_j^r) - (\pi_i^u - \pi_i^r) > 0$.

To prove sufficiency, suppose that the expression in (2.19) is strictly positive. Is this compatible with $v_i < v_j$. We show by contradiction that this cannot be the case. For this purpose assume that the expression in (2.19) is > 0 and that $v_i < v_j$. Using the fact that $v_i < v_j$ we can rewrite the expression (2.19). Note that

$$\begin{aligned}
& E[\max\{x_i, x_j\}\,|v_j \le x_i \le v_i, v_j \le x_j \le v_i] \\
= \;& E[\max\{x_i, x_j\}\,|v_i \le x_i \le v_j, v_i \le x_j \le v_j] \text{ for } v_i < v_j.
\end{aligned}$$

We therefore obtain from (2.19) by the assumption that $v_i < v_j$:

$$\begin{aligned}
&G_j(v_j)(G_i(v_i) - G_i(v_j))E[x_i\,|v_i \le x_i \le v_j]] + G_j(v_j)[G_i(v_i)v_i - G_i(v_j)v_j] \\
&+(G_j(v_i) - G_j(v_j))G_i(v_j)[v_i + E[x_j\,|v_i \le x_j \le v_j]] \qquad (2.20)\\
&+(G_j(v_i) - G_j(v_j))(G_i(v_i) - G_i(v_j))[v_i - E[\max\{x_i, x_j\}\,|v_i \le x_i \le v_j, v_i \le x_j \le v_j]].
\end{aligned}$$

It is easy to see that (2.20) is strictly negative when $v_i < v_j$, which constitutes a contradiction to the premise that the expression in (2.19) is > 0. Hence, if the latter is true, it must hold that $v_i > v_j$.

Q.E.D.

Proof of Proposition 2.6: In the following, we derive the restrictions that incentive compatibility of the inspection procedure depicted in Figure 2.4 imposes on the wage-schedule given in (2.8). Where appropriate, we will simplify the components of an incentive compatible wage-schedule in line with the principal's goal of minimizing the (expected) payment to the agent.

(i) **Ex Post Stage**: Suppose the agent has inspected alternatives i and j and has found payoffs x_i and x_j, respectively. If he submits reports $(\widehat{x}_i, \widehat{x}_j)$, where $\widehat{x}_i > \widehat{x}_j$, the principal will choose alternative i, in which case the agent receives the wage $w_1(x_i)$ if the report is correct, and is punished otherwise. Likewise, if the agent submits reports $(\widehat{x}_i, \widehat{x}_j)$, where $\widehat{x}_i < \widehat{x}_j$, the principal chooses alternative j. In this case, the agent receives the wage $w_2(x_i)$ if the report is correct, and is punished otherwise. As the agent can misrepresent the payoff associated with the alternative *not* chosen by the principal, ex post incentive compatibility requires that $w_1(x_i) \geq w_2(x_j)$ whenever alternative i's payoff x_i is greater than alternative j's payoff x_j. Likewise, it must hold that $w_1(x_i) \leq w_2(x_j)$ whenever alternative j's payoff is greater than alternative i's. Suppose now that the agent has found identical payoff realizations for the two alternatives (i.e. $x_i = x_j = x$, where $x \in (v_j, \overline{x}_i]$). If it holds under an incentive compatible wage-schedule that $w_1(x) \neq w_2(x)$, then agent will find it optimal to ensure that the principal chooses the alternative which maximizes the agent's wage. He can do so by understating the payoff associated with the alternative for which the wage is lower. Incentive compatibility therefore requires

$$w_2(x) = w_1(x) \text{ for all } x \in (v_j, \overline{x}_i],$$

where w_1 is non-decreasing on $(v_j, \overline{x}_i]$ to satisfy ex post incentive compatibility.

(ii) **Interim Stage**: Suppose the agent has inspected alternative i and has found a payoff realization x_i for which valuable information can be gained from an inspection of alternative j (i.e. $x_i \in (v_j, \overline{x}_i]$). The agent will inspect alternative j if and only if

$$w_1(x_i) \leq w_1(x_i)G_j(x_i) + \int_{x_i}^{\overline{x}_j} w_2(x_j)dG_j(x_j) - c.$$

Re-arranging terms and using the fact that $w_1(x) = w_2(x)$ for all $x \in (v_j, \overline{x}_i]$, we can write this interim incentive constraint as follows:

$$\begin{aligned} c \;\leq\; & (G_j(\overline{x}_i) - G_j(x_i)) \left(E[w_1(x_j) \,|x_i < x_j \leq \overline{x}_i] - w_1(x_i)\right) \\ & + (1 - G_j(\overline{x}_i)) \left(E[w_2(x_j) \,|x_j > \overline{x}_i] - w_1(x_i)\right) \end{aligned} \tag{2.21}$$

The right hand side of the inequality in (2.21) represents the expected wage increase that the agent can get by inspecting alternative j. An optimal contract will make this expected increase as small as possible. As ex post incentive compatibility requires the function w_1 to be non-decreasing, it is optimal to set $w_1(x) = \underline{w}$ for all $x \in (v_j, \overline{x}_i]$, thus providing the agent (weak) ex post incentives to truthfully reveal both payoff realizations correctly. Note that making w_1 constant is without loss of generality for ex ante incentive provision since, at the ex ante stage, the agent's incentives are determined only by the expectation of w_1. The interim incentive constraint in (2.21) therefore reduces to

$$E[w_2(x_j)\,|x_j > \overline{x}_i] \geq \underline{w} + \frac{c}{1 - G_j(\overline{x}_i)}. \qquad \text{(Interim IC)}$$

It is optimal for the principal to make this interim incentive constraint bind in order to extract the agent's entire surplus from the second inspection. Without loss of generality we can set $w_2(x_j) = \underline{w} + \frac{c}{1-G_j(\overline{x}_i)}$ for all $x_j > \overline{x}_i$.[17] Thus, an optimal incentive compatible wage-schedule of the form (2.8) reduces to

$$w = \begin{cases} w_0(x_j) & \text{if } x_i \leq v_j \\ \underline{w} & \text{if } x_i > v_j \text{ and } x_j \leq \overline{x}_i \\ \underline{w} + \frac{c}{1-G_j(\overline{x}_i)} & \text{if } x_i > v_j \text{ and } x_j > \overline{x}_i \end{cases} \qquad (2.22)$$

(iii) **Ex ante Stage**: To ensure ex ante incentive compatibility it must hold that no deviation at the ex ante stage yields the agent a higher expected payoff than that associated with the desired inspection depicted in Figure 2.4, which is given by

$$G_i(v_j)E[w_0(x_j)] + (1 - G_i(v_j))\underline{w} - c =: \Pi^* \qquad (2.23)$$

First suppose that the agent does not inspect any alternatives at all. In order to avoid the risk of punishment, he must submit any report for which the principal chooses alternative j immediately (i.e. $\widehat{x}_i \leq v_j$). The agent's expected wage from this procedure is $E[w_0(x_j)]$. To deter the agent from adopting this procedure it

[17] All that matters for incentive provision at the interim stage is the conditional expectation of w_2 on $(\overline{x}_i, \overline{x}_j]$.

must hold that $\Pi^* \geq E[w_0(x_j)]$, which implies that

$$\underline{w} \geq E[w_0(x_j)] + \frac{c}{1 - G_i(v_j)}. \tag{2.24}$$

Thus, $\underline{w} > E[w_0(x_j)]$, it hence it follows immediately that $\Pi^* < \underline{w}$.

Now consider the procedure under which the agent first inspects alternative j. If he finds that $x_j > v_j$ he can pretend that he has inspected both alternatives, but that j is better. If instead $x_j \leq v_j$ he retrospectively inspects alternative i. This deviation yields the following expected payoff

$$G_j(v_j)[G_i(v_j)E[w_0(x_j)\,|x_j \leq v_j] + (1 - G_i(v_j))\underline{w} - c] + (1 - G_j(v_j))\underline{w} =: \widehat{\Pi} \tag{2.25}$$

Using the expression for Π^* in (2.23) we can write (2.25) as follows

$$\widehat{\Pi} = G_j(v_j)\left(\Pi^* - G_i(v_j)\left(E[w_0(x_j)] - E[w_0(x_j)\,|x_j \leq v_j]\right)\right) + (1 - G_j(v_j))\underline{w}. \tag{2.26}$$

Ex ante incentive compatibility of the wage-schedule in (2.22) requires that $\Pi^* \geq \widehat{\Pi}$, which, in turn, implies that

$$\Pi^* \geq \underline{w} - \frac{G_j(v_j)G_i(v_j)}{(1 - G_j(v_j))}\left(E[w_0(x_j)] - E[w_0(x_j)\,|x_j \leq v_j]\right).$$

From above we know that $\Pi^* < \underline{w}$, and therefore a *necessary* condition for ex ante incentive compatibility is that

$$E[w_0(x_j)] > E[w_0(x_j)\,|x_j \leq v_j]. \tag{2.27}$$

Condition (2.27) places mild restrictions on the shape of the function w_0 - it is easy to see that w_0 cannot be constant for all x_j. Condition (2.27) roughly states that w_0 must reward the agent more for realizations above v_j than for realizations below v_j. The exact functional form of w_0, however, is irrelevant for ex ante incentive provision, as only the expectation of w_0 matters. Thus, it is sufficient to consider step functions for which $w_0(\cdot)$ takes on a fixed value for all $x_j \leq v_j$, and a strictly higher value for at least some realizations $x_j > v_j$. There is no loss of optimality in restricting ourselves to step functions of the following form:[18]

$$w_0(x_j) = \begin{cases} 0 & \text{if } x_j \leq \overline{x}_i \\ \overline{w} & \text{if } x_j > \overline{x}_i \end{cases} . \tag{2.28}$$

Thus, $E[w_0(x_j)] = \overline{w}\,(1 - G_j(\overline{x}_i))$ and $E[w_0(x_j)\,|x_j \leq v_j] = 0$. An optimal incentive compatible wage-schedule therefore reduces to

$$w = \begin{cases} 0 & \text{if } x_i \leq v_j \text{ and } x_j \leq \overline{x}_i \\ \underline{w} & \text{if } x_i > v_j \text{ and } x_j \leq \overline{x}_i \\ \overline{w} & \text{if } x_i \leq v_j \text{ and } x_j > \overline{x}_i \\ \underline{w} + \frac{c}{1-G_j(\overline{x}_i)} & \text{if } x_i > v_j \text{ and } x_j > \overline{x}_i \end{cases} \tag{2.29}$$

The optimal values of $\underline{w}$ and $\overline{w}$ are determined by the two ex ante incentive constraints described above. Using the step function in (2.28), the ex ante incentive constraint in (2.24) becomes

$$\underline{w} \geq (1 - G_j(\overline{x}_i))\,\overline{w} + \frac{c}{(1 - G_i(v_j))}. \qquad \text{(Ex Ante IC 1)}$$

By (2.28) the second ex ante incentive constraint (namely that $\Pi^* \geq \widehat{\Pi}$) reduces to

$$\overline{w} \geq \frac{(1 - G_j(v_j))\,c}{G_i(v_j)\,(1 - G_j(\overline{x}_i))} + \frac{(1 - G_j(v_j))}{(1 - G_j(\overline{x}_i))}\underline{w}. \qquad \text{(Ex Ante IC 2)}$$

It is optimal for the principal to make the two ex ante incentive constraints bind, so that the optimal values of $\underline{w}$ and $\overline{w}$ are characterized by the solution to a system of two linear equations; in particular,

$$\underline{w} = \frac{(G_i(v_j) + (1 - G_i(v_j))(1 - G_j(v_j)))\,c}{G_j(v_j)G_i(v_j)(1 - G_i(v_j))}$$

and

$$\overline{w} = \frac{(1 - G_j(v_j))c}{G_i(v_j)(1 - G_j(\overline{x}_i))} + \frac{(1 - G_j(v_j))}{(1 - G_j(\overline{x}_i))}\underline{w}.$$

Q.E.D.

[18]Recall that $\overline{x}_i > v_j$.

Chapter 3

How much do you really care? A study of Efficient Compromising[1]

3.1 Introduction

You and your partner disagree about which restaurant to visit tonight. You prefer the Italian restaurant over the English restaurant, and the English restaurant over the Chinese restaurant. But your partner has exactly the opposite preferences. Should you compromise, and go to the English restaurant, or should you go to a restaurant that one of you likes best? The answer presumably depends on how much each partner really cares whether his or her most preferred alternative or the compromise is chosen. Is there a way of finding out how much each partner cares, or will they always necessarily overstate the importance of the decision? This is the question which this chapter addresses.

How we answer our question depends, of course, on what precisely we mean by "strength" of preference. This could be interpreted as the amount of money that an agent is willing to pay to obtain one outcome rather than another. If this is what we have in mind when speaking of "strength" of preference, then one could try to elicit the strength of the partners preferences by introducing a mechanism according to which a

[1]This chapter is joint work with Tilman Börgers.

partner whose preferred restaurant is chosen has to pay the other some compensation.

Here, we abstract from such side payments, as they seem inappropriate in many situations. Spouses, for example, rarely pay money to each other to resolve conflicts. When initially conceiving of this project, we had another situation in mind in which money payments are typically not made: voting. Optimal voting rules, if there are more than two candidates, need to elicit, in some sense, the "strength" of preference for candidates, yet voters are typically not asked to offer payments in conjunction with their votes. The problem that we study here is a simplified version of the voting problem.

If side payments are ruled out, what do we mean by "strength" of preferences, and how can we elicit them? We mean in this chapter by strength of preference the von Neumann Morgenstern utility of different alternatives. If we evaluate different mechanisms from an *ex ante* perspective (Holmström and Myerson (1983)) then von Neumann Morgenstern utilities have to be taken into account when resolving conflicts. How can we hope to elicit von Neumann Morgenstern utilities truthfully? By exposing agents to risk. Agents' choices among lotteries indicate their von Neumann Morgenstern utilities. If agents play a game with incomplete information, then they are almost always automatically exposed to risk. Their choices can then help with efficient decision making.

We develop this theme in a simple stylized example with two agents and three alternatives. We assume that it is known that the agents' ranking of the alternatives is diametrically opposed. Their von Neumann Morgenstern utilities for the alternatives are, however, not known. It is optimal to implement the compromise if and only if the sum of the agents' utilities of the compromise is larger than the sum of the agents' utilities from their most preferred alternatives. A decision rule with this property is called *first best.*

In simple examples, in which there are only two or three values of the von Neumann Morgenstern utility of the compromise to which the prior distribution attaches positive values, we find that simple mechanisms implement the first best decision rule. However, for the case that a continuum of values of the von Neumann Morgenstern utilities of the compromise is in the support of the prior distribution, we prove an impossibility: There is no incentive compatible decision rule that implements the compromise efficiently.

The proof of the impossibility result in the present work is related to the proof of the impossibility of efficient bilateral trade due to Myerson and Satterthwaite (1983). Our argument uses the incentive compatibility constraints to obtain a differential equation for the probabilities with which different alternatives are implemented by an incentive compatible first best rule. This is similar to Myerson and Satterthwaite's use of a differential equation to derive the payments made by the agents in an incentive compatible first best trading rule. Where Myerson and Satterthwaite then show that an ex ante budget constraint is violated by these payments, we show that the constraint that probabilities need to add up to one is violated ex ante.

The present work is also remotely related to Börgers (1991), where the impossibility to compromise efficiently was demonstrated in a setting in which lotteries played no role, but agents' information about other agents' preferences was assumed to be much poorer than it is in the present work.

Finally, the present chapter is related to a recent paper by Li, Rosen and Suen (2001), who study incentive compatible decision rules in a common value setting: committee members have private information about the underlying state of the world which affects their payoffs from the collective decision. Li et al. show that incentive compatible binary deterministic decision rules partition agents' private information, where the degree of conflict between the agents' interests determines the coarseness of the partition.

This chapter is organized as follows. In Section 3.2 we introduce our model, extensions to which are briefly discussed in Section 3.3. Section 3.4 describes normative properties of decision rules. In particular, we explain what we understand by "first best" decision rules. Section 3.5 defines incentive compatibility in our setting, and obtains some general implications of incentive compatibility. In Sections 3.6 and 3.7 we then deal first with the case that there are only two possible values of the von

Neumann Morgenstern utility of the compromise, and later with the case that there are three such values. The purpose of these sections is to show in simple examples that it is possible to elicit von Neumann Morgenstern utilities through simple mechanisms, and that sometimes it is even possible to do so and to implement the first best. Section 3.8 then deals with the case in which there is a continuum of possible values of the von Neumann Morgenstern utility of the compromise. We prove a rather general impossibility result. Section 3.9 concludes.

3.2 Model

There are two agents: $i = 1, 2$ who must collectively choose one alternative from the set $\{A, B, C\}$. Agent 1 prefers A over B, and B over C. Agent 2 prefers C over B, and B over A. These preferences are common knowledge among the two agents. Each agent i has a von Neumann Morgenstern utility function: $u_i : \{A, B, C\} \to \mathbb{R}$. We normalize each agent's utility from her least preferred alternative to zero: $u_1(C) = u_2(A) = 0$. In addition, we assume that each agent's utility from her most preferred alternative is 1: $u_1(A) = u_2(C) = 1$. This assumption means that agent 1's *intensity of preference* for her most preferred alternative over her least preferred alternative is the same as agent 2's.[2] These features of the von Neumann Morgenstern utility functions are common knowledge among the two agents.

For $i = 1, 2$ we write t_i for $u_i(B)$. We refer to t_i as "player i's type". We assume that t_i is a random variable which is only observed by agent i. The two players' types are stochastically independent, and they are identically distributed with common cumulative distribution function G. We assume that G has support $S \subseteq [0, 1]$. We consider only the two special cases: (i) the case that S is finite, and (ii) the case that $S = [0, 1]$ and that G has a strictly positive density on $[0, 1]$. The joint distribution of (t_1, t_2) is common knowledge among the agents.

A decision rule f is a function $f : S^2 \to \Delta(\{A, B, C\})$ where $\Delta(\{A, B, C\})$ is the set of probability distributions over $\{A, B, C\}$. We write $f_A(t_1, t_2)$ for the probability which

[2] For a brief discussion of this assumption, see Section 3.3 below.

$f(t_1, t_2)$ assigns to alternative A, and we define $f_B(t_1, t_2)$ and $f_C(t_1, t_2)$ analogously. Given any decision rule, denote for every $t_i \in S$ by $p_i(t_i)$ the conditional probability that the alternative that agent i likes best is implemented, conditional on agent i's type being t_i, i.e.:

$$p_1(t_1) = \int_S f_A(t_1, t_2) dG(t_2) \quad \text{and} \quad p_2(t_1) = \int_S f_C(t_1, t_2) dG(t_1).$$

Denote by $q_i(t_i)$ the probability that the compromise is implemented, conditional on agent i's type being t_i, i.e. for $i = 1, 2$:

$$q_i(t_i) = \int_S f_B(t_1, t_2) dG(t_j) \quad \text{where} \quad j \neq i.$$

Finally, we denote by $U_i(t_i)$ agent i's expected utility, conditional on being type t_i, that is:

$$U_i(t_i) = p_i(t_i) + q_i(t_i) t_i.$$

3.3 Discussion on Intensity of Preference

The special feature of the model described in Section 3.2 is that both agents have identical intensity of preference for their most preferred alternative over their least preferred alternative. In general, there is no reason to believe that two agents' cardinal von Neumann Morgenstern utilities should satisfy this assumption, even in situations where the agents' ordinal preferences over alternatives are known to be diametrically opposed. We can, however, view the model described in Section 3.2 as the "reduced form" of a more general setup in which agents' intensities of preference differ.

To see this, let agent i's von Neumann Morgenstern utility function be given by the numbers $u_i(A)$, $u_i(B)$ and $u_i(C)$. As before, normalize each agent's utility from her least-preferred alternative to zero (i.e. $u_1(C) = u_2(A) = 0$). Now, however, suppose that each agent's value of her most-preferred alternative is believed to be the realization of a random variable ρ with some known distribution on $\mathbb{R}_+$. That is, $u_1(A) \equiv \rho_1$ and $u_2(C) \equiv \rho_2$, where $\rho_1, \rho_2 \in \mathbb{R}_+$ are drawn independently of one another. Note that ρ_i measures agent i's intensity of preference for her most preferred over her least preferred alternative. Finally, assume that agent i's utility $u_i(B)$ from the compromise is given by $t_i\rho_i$, where t_i is drawn (independently of ρ_i, ρ_j and t_j) from G on $[0,1]$.

In this setup, agent i's private information is given by $(\rho_i, t_i) \in \mathbb{R}_+ \times [0,1]$. As above, a collective decision rule prescribes a (possibly randomized) choice from the set of alternatives $\{A, B, C\}$ on the basis of the agents' private information.[3] For concreteness assume furthermore that agents' preference intensity can only take on one of two values: $\underline{\rho}, \overline{\rho}$, where $\overline{\rho} > \underline{\rho}$. Now take a decision rule that gives each agent incentives to truthfully reveal her private information. Consider agent 1 and suppose that if she reports her private information as $(\underline{\rho}, t)$, the decision rule selects the lottery $l \in \Delta(\{A, B, C\})$.[4] If, instead, she reports her private information as $(\overline{\rho}, t)$, the decision rule selects the lottery $l' \in \Delta(\{A, B, C\})$, where $l' \neq l$.

Now suppose that agent 1 has "preference intensity" $\underline{\rho}$. As the decision rule is assumed to be incentive compatible, agent 1 prefers to submit the report $(\underline{\rho}, t)$ over the report $(\overline{\rho}, t)$. In the former case, her expected utility is $\underline{\rho}(l_A + l_B t)$, while in the latter it is $\underline{\rho}(l'_A + l'_B t)$. Hence, incentive compatibility implies

$$\underline{\rho}(l_A + l_B t) \geq \underline{\rho}(l'_A + l'_B t). \tag{IC$_{\underline{\rho}}$}$$

Now suppose that agent 1 has "preference intensity" $\overline{\rho}$. As the decision rule is assumed to be incentive compatible, agent 1 prefers to submit the report $(\overline{\rho}, t)$ over the report $(\underline{\rho}, t)$. In the former case, her expected utility is $\overline{\rho}(l'_A + l'_B t)$, while in the latter it is $\overline{\rho}(l_A + l_B t)$. This implies that

$$\overline{\rho}(l'_A + l'_B t) \geq \overline{\rho}(l_A + l_B t). \tag{IC$_{\overline{\rho}}$}$$

Combining IC$_{\underline{\rho}}$ and IC$_{\overline{\rho}}$ implies that

$$l'_A + l'_B t = l_A + l_B t.$$

[3]Formally, a decision rule now is a mapping $\varphi : (\mathbb{R}_+ \times [0,1])^2 \to \Delta(\{A, B, C\})$ where $\varphi_k(\rho_1, \rho_2, t_1, t_2)$ denotes the probability that alternative $k \in \{A, B, C\}$ is chosen when agent 1's information is (ρ_1, t_1) and agent 2's information is (ρ_2, t_2).

[4]Note that $l = (l_A, l_B, l_C)$, where $l_C = 1 - l_A - l_B$.

Thus, if agent 1 claims her private information is $(\underline{\rho}, t)$, then she has the same expected utility as if she had reported it as $(\overline{\rho}, t)$. Therefore, a decision rule that assigns the lottery l to both reports $(\underline{\rho}, t)$ and $(\overline{\rho}, t)$ is also incentive compatible. In this sense, we can say that any incentive compatible decision rule must ignore agents' intensity of preference for their most over their least preferred alternatives. We may therefore restrict ourselves to the model presented in Section 3.2 in which the intensity of preference only takes a single value (normalized to 1).

3.4 Normative Properties of Decision Rules

We calculate the expected welfare resulting from a decision rule f using a utilitarian welfare criterion. This corresponds to the evaluation of the ex ante expected utility of an agent who does not know whether he will be agent 1 or 2, and who does not yet know his type. We assume that the probability of being either agent 1 or agent 2 is equal to 1/2.

Definition 3.1 *The ex ante expected utility associated with decision rule f is:*

$$\int_S U_1(t_1)dG(t_1) + \int_S U_2(t_2)dG(t_2).$$

It is easy to see that the decision rules f that maximize ex ante expected utility among all decision rules are those that are "first best" in the sense of the following definition.

Definition 3.2 *A decision rule f is called "first best" if*

$$\begin{aligned} t_1 + t_2 &> 1 \Rightarrow f_B(t_1, t_2) = 1 and \\ t_1 + t_2 &< 1 \Rightarrow f_B(t_1, t_2) = 0. \end{aligned}$$

Note that there are many first best decision rules. The reason is that the above definition does not restrict in any way the probabilities with which A and C are chosen if the compromise is *not* implemented. A second normative property that will play a role in this chapter is the following symmetry condition.

Definition 3.3 *A decision rule f is called "symmetric" if for all $t_1, t_2 \in S$ we have:*

$$f_A(t_1, t_2) = f_C(t_2, t_1).$$

Lemma 3.1: *If a decision rule f is symmetric then for all $t_1, t_2 \in S$ we have:*

$$f_B(t_1, t_2) = f_B(t_2, t_1).$$

Proof of Lemma 3.1: By definition, $f_C(t_2, t_1) = 1 - f_A(t_2, t_1) - f_B(t_2, t_1)$. Using Definition 3.3 to replace $f_A(t_2, t_1)$ we obtain

$$\begin{aligned} f_C(t_2, t_1) &= 1 - f_C(t_1, t_2) - f_B(t_2, t_1) \\ &= 1 - (1 - f_A(t_1, t_2) - f_B(t_1, t_2)) - f_B(t_2, t_1) \\ &= f_A(t_1, t_2) + f_B(t_1, t_2) - f_B(t_2, t_1). \end{aligned}$$

Applying Definition 3.3 once more yields $f_B(t_2, t_1) = f_B(t_1, t_2)$.

Q.E.D.

3.5 Incentive Compatibility

Because types are privately observed, a decision rule can be implemented in practice only if it is incentive compatible.[5]

Definition 3.4 *A decision rule f is "incentive compatible" if for $i = 1, 2$ and for any types $t_i, t_i' \in S$ we have:*

$$p_i(t_i) + q_i(t_i)t_i \geq p_i(t_i') + q_i(t_i')t_i.$$

The following simple Lemma is key for understanding how an incentive compatible rule incentivizes agents to reveal their true von Neumann Morgenstern utility of the compromise.

[5]The following definition implicitly assumes that the mechanism which is used to implement the decision rule is a direct one. The "Revelation Principle" (see, for example, Proposition 5.1 in Krishna (2002)) says that this assumption is without loss of generality.

Lemma 3.2 *A decision rule f is "incentive compatible" if and only if for $i = 1, 2$ we have:*

(i) *q_i is (weakly) monotonically increasing in t_i;*

(ii) *p_i is (weakly) monotonically decreasing in t_i, and $p_i(t_i) \neq p_i(t_i') \Leftrightarrow q_i(t_i) \neq q_i(t_i')$;*

(iii) *for any two types $t_i, t_i' \in S$ with $t_i < t_i'$, if $q_i(t_i) \neq q_i(t_i')$ we have:*

$$t_i \leq \frac{p_i(t_i) - p_i(t_i')}{q_i(t_i') - q_i(t_i)} \leq t_i'.$$

In words, Lemma 3.2 says that the probability of the compromise, conditional on an agent's type, increases as this agent's utility of the compromise increases (item (i)). Where is this probability taken from? If it were taken from the agent's least preferred alternative only, then the agent would have an incentive to pretend to have a higher utility for the compromise than he actually has. Therefore, the probability must be taken in parts from the alternative which the agent likes best (item (ii)). The agent "pays" for a higher probability of the compromise with a lower probability of his most preferred alternative.

Item (iii) of the lemma concerns the ratio of the reduction in the probability of the most preferred alternative and the increase in the probability of the compromise. In other words this is the price paid per unit increase in the probability of the compromise. Item (iii) says that this price needs to be increasing in the agent's type, that is, agents who value the compromise more pay a higher price than agents who value it less. This is how an incentive compatible decision rule screens low type agents and high type agents.

Proof of Lemma 3.2: The condition that defines incentive compatibility can equivalently be written with reference to all pairs of types $t_i, t_i' \in S$ which are ordered, i.e. $t_i < t_i'$. We then need to require:

$$p_i(t_i) + q_i(t_i)t_i \geq p_i(t_i') + q_i(t_i')t_i \tag{3.1}$$

$$p_i(t_i) + q_i(t_i)t_i' \leq p_i(t_i') + q_i(t_i')t_i' \tag{3.2}$$

Subtracting (3.2) from (3.1) we obtain:

$$q_i(t_i)(t_i - t_i') \geq q_i(t_i')(t_i - t_i')$$

which shows that (i) in Lemma 3.2 is necessary. The necessity of (ii) follows from the argument given in words in the paragraph preceding this proof. Finally (iii) is obtained simply by rewriting (3.1) and (3.2).

To see the sufficiency of (i)-(iii) note that (i) and (ii) imply that for any pair of types, $t_i, t_i' \in S$, we either have that $q_i(t_i) = q_i(t_i')$ and $p_i(t_i) = p_i(t_i')$, in which case it is obvious that neither type has an incentive to pretend to be the other type, or we have that q_i is strictly larger and p_i is strictly lower for the higher type. But in that case the inequality in (iii) implies that the incentive compatibility constraints hold.

Q.E.D.

We now focus on the special case that G has a density that is strictly positive everywhere on $[0, 1]$. The next Lemma shows how some standard results from mechanism design with transferable utility generalize to our context.[6]

Lemma 3.3 *Suppose that G has a density that is positive everywhere on $[0, 1]$. If a decision rule f is incentive compatible, then for every agent $i = 1, 2$:*

(i) *U_i is differentiable almost everywhere, and in all points in which U_i is differentiable:*

$$U_i'(t_i) = q_i(t_i);$$

(ii) *for every $t_i \in [0, 1]$:*

$$U_i(t_i) = U_i(1) - \int_{t_i}^{1} q_i(s_i) ds_i.$$

In words, point (i) and the fact that for incentive compatible mechanisms q_i is increasing in t_i (part (i) of Lemma 3.2) imply that U_i is convex in t_i. Part (ii) shows

[6] See pages 64 and 65 in Krishna (2002).

that the expected utility which an agent derives from the mechanism is completely determined by the expected utility $U_i(1)$ of the highest type, and by the expected probabilities of the compromise, conditional on the agent's type. This is analogous to the principle in auction theory that the expected utility of a bidder is determined up to a constant by the allocation rule that determines the probability with which the bidder obtains the object.

Proof of Lemma 3.3: Incentive compatibility implies that

$$U_i(t_i) = \max_{t_i' \in [0,1]} \{p_i(t_i') + q_i(t_i')t_i\}.$$

Thus, U_i is the maximum of a set of affine functions and is therefore convex. This implies that U_i is differentiable almost everywhere in $(0, 1)$. It follows from the Envelope Theorem that, at every point at which U_i is differentiable,

$$U_i'(t_i) = q_i(t_i). \tag{3.3}$$

As U_i is the definite integral of its derivative, we have

$$U_i(t_i) = U_i(1) - \int_{t_i}^{1} q_i(s_i)ds_i.$$

This implies that up to the constant U_i, agent i's expected utility from any incentive compatible decision rule depends *only* on the probability of the compromise. By definition, $U_i(t_i) = p_i(t_i) + q_i(t_i)t_i$, so that

$$p_i(t_i) = U_i(1) - q_i(t_i)t_i - \int_{t_i}^{1} q_i(s_i)ds_i, \tag{3.4}$$

where $U_i(1) = p_i(1) + q_i(1)$.

Q.E.D.

Our interest will be in those decision rules that maximize ex ante expected utility among all incentive compatible rules. We make the following

Definition 3.5 *A decision rule f is called "second best" if it yields the largest ex ante expected utility among all incentive compatible decision rules.*

The following simple result shows that when considering whether first best rules are incentive compatible, or what second best rules would be, as we do in this chapter, there is no loss of generality in considering symmetric decision rules only.

Lemma 3.4 *For every incentive compatible decision rule f there is a symmetric incentive compatible decision rule f' which yields the same ex ante expected utility as f.*

To see why this is true define f' to be the rule which results if agents are first asked to reveal their types, then a fair coin is tossed and, if head comes up f is applied, but if tails comes up f is applied except that the roles of the agents are reversed, i.e. agent 2 now plays the role of agent 1, and vice versa. Because f is incentive compatible, both agents would have no incentive to distort their preferences if they knew the outcome of the coin toss in advance. Therefore they also have no incentive to distort preferences ex ante. It is obvious that the new rule f' is symmetric, and that it has the same ex ante expected utility as f.

3.6 Two types

We ask ourselves whether first best rules are incentive compatible, and, if they are not, what second best incentive compatible rules would be. We begin with a look at the special case that $S = \{t, t'\}$ where $t < t\prime$. If both types are smaller than 0.5, or both types are larger than 0.5, then there are symmetric first best rule that are constant. In the former case such a rule is: always implement the lottery $(0.5, 0, 0.5)$ (listing first the probability of A, then the probability of B, and then the probability of C)[7]. In the latter case, such a rule is: always implement the lottery $(0, 1, 0)$. These rules are obviously incentive compatible.

Consider the case that $t < 0.5 < t'$. For simplicity we ignore the situation that one of the two types equals 0.5. We need to distinguish between the case in which $t + t' < 1$, and the case that $t + t' > 1$. For simplicity we ignore the case that $t + t' = 1$.

[7] We shall use the same notation throughout the paper.

In the former case, an example of a symmetric first best rule is in Figure 3.1, where rows refer to agent 1's type, and columns refer to agent 2's type.

1\2	t	t'
t	(0.5, 0, 0.5)	(0.5, 0, 0.5)
t'	(0.5, 0, 0.5)	(0, 1, 0)

Figure 3.1: Unanimity required to implement Compromise

Note that in Figure 3.1 it is weakly dominant for each agent to tell his true type. Thus, this first best rule is incentive compatible.

If $t + t' > 1$, then the following decision rule is first best and incentive compatible:

1\2	t	t'
t	(0.5, 0, 0.5)	(0, 1, 0)
t'	(0, 1, 0)	(0, 1, 0)

Figure 3.2: Unanimity required to veto Compromise

The two mechanisms in Figures 3.1 and 3.2 reflects mechanisms where, in the case of Figure 3.1, unanimity is required to implement the compromise and in the case of Figure 3.2, unanimity is required to veto the compromise. We conclude that in the case of two types the first best is incentive compatible. It can even be implemented in dominant strategies.

It is worthwhile at this stage to briefly consider an alternative (indirect) mechanism that has been advocated in practice: *approval voting*. Approval voting has been designed for much more general environments than we consider here (see Brams and Fishburn, 1983). Under approval voting each agent indicates for each alternative whether he approves of it, or not. The alternative that is approved of by the largest number of agents is then chosen. If several alternatives receive the same number of votes, then each of them is picked with the same probability.

It is easy to see what approval voting implies in our context. In the unique equilibrium of approval voting agents with type below 0.5 vote for their preferred alternative only. Agents with type above 0.5 vote for their most preferred alternative, and the compromise. Agents whose type equals exactly 0.5 are indifferent between the two strategies. Assuming as we did above that $t < 0.5 < t'$, Figure 3.3 shows the decision rule that is implemented by approval voting.

1\2	t	t'
t	$(\frac{1}{2}, 0, \frac{1}{2})$	$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
t'	$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$	$(0, 1, 0)$

Figure 3.3: Decision Rule implemented by Approval Voting

Now this decision rule is first best only if $t + t' = 1$, a non-generic case ruled out above. The background to this is that approval voting implements a decision rule that sometimes randomizes between the compromise and the agents' most preferred alternatives. In the first best, such a randomization occurs only in the non-generic case that $t + t' = 1$.

3.7 Three types

We now consider the case of three types: $S = \{t, t', t''\}$ where $t < t' < t''$. We denote the probabilities of the three types by $g(t)$, $g(t')$ and $g(t'')$. To identify the values of the exogenous parameters for which interesting complications arise, we begin by considering the probability of the compromise B under the first best rule. Notice that this probability is unique and is either zero or one, except in the non-generic case that the sum of two types is exactly 1. We rule out this case. Figure 3.4 shows the eight different forms that the function which assigns to each pair of types the first best probability of the compromise can take. Eight different cases arise.[8]

[8]Rows in Figure 3.4 refer to agent 1's type, columns to agent 2's.

	t	t'	t''
t	0	0	0
t'	0	0	0
t''	0	0	0

Case 1

	t	t'	t''
t	0	0	0
t'	0	0	0
t''	0	0	1

Case 2

	t	t'	t''
t	0	0	0
t'	0	0	1
t''	0	1	1

Case 3

	t	t'	t''
t	0	0	0
t'	0	1	1
t''	0	1	1

Case 4

	t	t'	t''
t	0	0	1
t'	0	0	1
t''	1	1	1

Case 5

	t	t'	t''
t	0	0	1
t'	0	1	1
t''	1	1	1

Case 6

	t	t'	t''
t	0	1	1
t'	1	1	1
t''	1	1	1

Case 7

	t	t'	t''
t	1	1	1
t'	1	1	1
t''	1	1	1

Case 8

Figure 3.4: FB probability of Compromise

Figure 3.5 shows for all possible values of (t, t') which case they give rise to. In Figure 3.5 the value of t'' is taken as given and fixed. The remaining two types, t and t' have to be smaller than t''. It also has to be the case that $t < t'$. The restrictions resulting from this are indicated in Figure 3.5 by solid lines. The remaining restrictions, corresponding to the various case distinctions, are indicated in dashed lines. Figure 3.5 assumes that $t'' > 0.5$. Thus, Case 1 of Figure 3.4 cannot arise in Figure 3.5. But all other cases are possible.

In Figure 3.4 it is obvious that the first best can be implemented in the trivial Cases 1 and 8. It is also easy to see that in Cases 2, 4, 5 and 7 the first best can be implemented, because in these cases the first best rule depends only on whether the players' types are "high" or "low", where one of the two categories, "high" or "low" contains two of the types. In all these cases one of the two (now indirect) mechanisms in Figures 3.1 and 3.2 will implement the first best.

Of interest are therefore only Cases 3 and 6. We shall show first that the first best

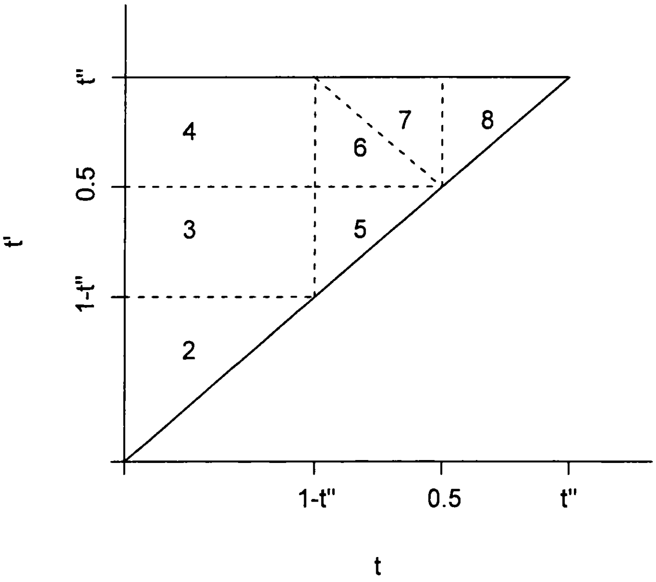

Figure 3.5: Parameter Regions (Numbers refer to Cases in Figure 3.4)

1\2	t	t'	t''
t	$(\frac{1}{2}, 0, \frac{1}{2})$	$(1-\alpha, 0, \alpha)$	$(1-\alpha, 0, \alpha)$
t'	$(\alpha, 0, 1-\alpha)$	$(\frac{1}{2}, 0, \frac{1}{2}))$	$(0, 1, 0)$
t''	$(\alpha, 0, 1-\alpha)$	$(0, 1, 0)$	$(0, 1, 0)$

Figure 3.6: Implementing FB in Case 3 (where $\alpha = \frac{1}{2} + g(t'')(\frac{1}{2} - \frac{t+t'}{2})$)

can always be implemented in Case 3. Indeed, a direct mechanism that implements the first best in Case 3 is displayed in Figure 3.6. It is easily verified that the rule displayed in Figure 3.6 is incentive compatible.

Because Figure 3.6 shows the first mechanism in this chapter in which non-trivial stochastic outcomes are used to separate agents of different types from each other, it is worthwhile to consider this mechanism in more detail. We can calculate the probabilities of different alternatives, conditional on the agents' types, as follows:

$$
\begin{aligned}
p_i(t) &= g(t)\alpha + g(t')\tfrac{1}{2} + g(t'')\tfrac{t+t'}{2} & \quad q_i(t) &= 0 \\
p_i(t') &= g(t)\alpha + g(t')\tfrac{1}{2} & \quad q_i(t') &= g(t'') \\
p_i(t'') &= g(t)\alpha & \quad q_i(t'') &= g(t') + g(t'')
\end{aligned}
$$

Thus, for example, when comparing the report t' to the report t, agents find that for reporting the higher type they pay the price that the most preferred alternative is chosen less frequently, with the loss in probability being $g(t'')\frac{t+t'}{2}$, but they gain that the compromise is chosen more frequently, with the increase in probability being $g(t'')$. The price paid in terms of loss of probability of the preferred alternative per unit of gained probability of the compromise is $\frac{t+t'}{2}$. Types t will not find it worthwhile to pay this price, but types t' will find it worthwhile. In this way, the mechanism separates these two types. Similarly, when comparing the report t'' to the report t', agents find that if they report the higher type the probability of their most preferred alternative decreases further, by $g(t')\frac{1}{2}$, but the probability of the compromise increases by $g(t')$. Thus, the price paid in lost probability of the most preferred alternative per unit of increased probability of the compromise is $\frac{1}{2}$. Types t' will not find this worthwhile (because $t' < 1/2$), but types t'' will find it a price worth paying (because $t'' > 1/2$).

It is a curious feature of the mechanism in Figure 3.6 that while the probability with which an agent's most preferred alternative is implemented is decreasing in that agent's type from the *interim* perspective, this is not true from the *ex ante* perspective. For example, if the other agent reports that his type is t, then the probability that an agent's most preferred alternative is implemented is actually increasing in this agent's type from $\frac{1}{2}$ to α.

We now turn to Case 6. In this case, a symmetric decision rule that is first best must take the form indicated in Figure 3.7.

It is immediate that type t' has no incentive to pretend to be type t'', and vice versa, if and only if $\beta \in [t', t'']$. Noting that this implies that $\beta > 0.5$ one can then see that the only incentive constraint that still needs to be taken care of is the constraint

1\2	t	t'	t''
t	$(\frac{1}{2},0,\frac{1}{2})$	$(1-\beta,0,\beta)$	$(0,1,0)$
t'	$(\beta,0,1-\beta)$	$(0,1,0))$	$(0,1,0)$
t''	$(0,1,0)$	$(0,1,0)$	$(0,1,0)$

Figure 3.7: Implementing FB in Case 6

that type t must not have an incentive to pretend to be type t'. This is the case if and only if:

$$g(t)\frac{1}{2}+g(t')(1-\beta)+g(t'')t \geq g(t)\beta+g(t')t+g(t'')t \Leftrightarrow \frac{\beta-\frac{1}{2}}{\frac{1}{2}-t} \leq \frac{g(t')}{g(t)+g(t')}.$$

We can pick a β from the interval $[t', t'']$ that satisfies this constraint if and only if the constraint is satisfied for $\beta = t'$, i.e.

$$\frac{t'-\frac{1}{2}}{\frac{1}{2}-t} \leq \frac{g(t')}{g(t)+g(t')}.$$

For given t'', there will always be values of t and t' for which this inequality is violated, and therefore the first best cannot be implemented. Figure 3.8 indicates the relevant region. The line that separates Region 6a from Region 6b has slope $-\frac{g(t')}{g(t)+g(t')}$.

We now consider two examples, one from parameter region 6a, and one from parameter region 6b. In both examples we assume that all types are equally likely: $g(t) = g(t') = g(t'') = \frac{1}{3}$. In both examples let $t = 0.3$ and $t'' = 0.8$. In Example 1 we choose $t' = 0.55$, and in Example 2 we choose $t' = 0.65$. In Example 1 the first best can be implemented as shown in Figure 3.9.

For Example 2 our earlier calculations show that we cannot implement the first best. We have determined the second best decision rule numerically.[9] It is displayed in Figure 3.10.

[9]Second best decision rules can be found as the solution to a linear programming problem. To solve Example 2 we used the implementation of the simplex algorithm that is available in the *boot*-package for the open source software R (R Development Core Team, 2003).

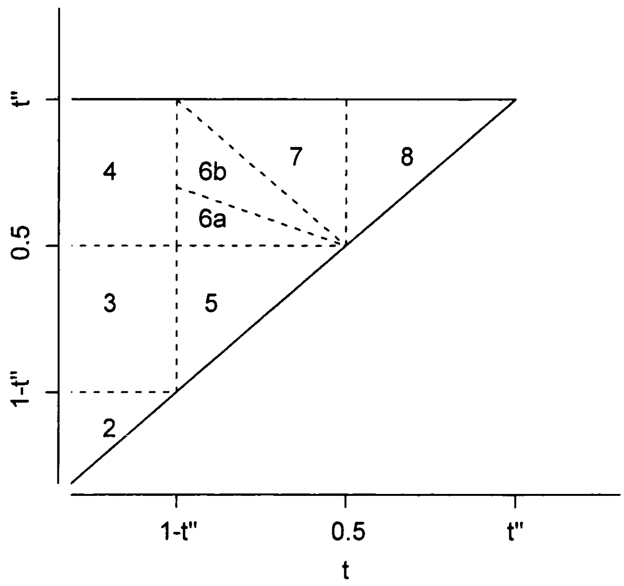

Figure 3.8: Region 6a: FB can be implemented; Region 6B: FB cannot be implemented

1\2	t	t'	t''
t	$(0.5, 0, 0.5)$	$(0.45, 0, 0.55)$	$(0, 1, 0)$
t'	$(0.55, 0, 0.45)$	$(0, 1, 0))$	$(0, 1, 0)$
t''	$(0, 1, 0)$	$(0, 1, 0)$	$(0, 1, 0)$

Figure 3.9: Implementing FB in Example 1

1\2	t	t'	t''
t	$(0.5, 0, 0.5)$	$(0.3825, 0, 0.6175)$	$(0.05, 0.95, 0)$
t'	$(0.6175, 0, 0.3825)$	$(0, 1, 0))$	$(0, 1, 0)$
t''	$(0, 0.95, 0.05)$	$(0, 1, 0)$	$(0, 1, 0)$

Figure 3.10: The Second Best in Example 2

The inefficiency occurs when one agent is of the lowest and the other agent is of the highest type. In this case, first best would require that the compromise is chosen with probability 1. But this cannot be made incentive compatible, and instead it is optimal to choose the compromise with probability 0.95 only. Changing this probability has a cascading effect: The intermediate type, when encountering the low type, can be rewarded with a lower probability of his or her most preferred alternative. Therefore, the low type has no incentive to pretend to be the intermediate type.

One can continue to explore second best rules for the parameter region 6b. The second best rule is sensitive to changes in the exogenous parameters. It does not appear that a "robust" second best mechanism emerges from this analysis.

3.8 Continuum of Types

We now turn to the case in which all types between 0 and 1 are in the support of the distribution G. For this case, we obtain an impossibility result.

Proposition 3.1 *Suppose that G has a density g that is strictly positive everywhere on $[0, 1]$. Then no first best decision rule is incentive compatible.*

Before we give the formal proof of this result, it may be worthwhile to outline the structure of our proof. The proof is indirect. We postulate the existence of a first best rule that is incentive compatible. We then first note that the probability with which the compromise is implemented, conditional on an agent's type, is determined by the definition of first best. The corresponding conditional probability with which any agent's most preferred alternative is implemented is then determined by the differential equation in Lemma 3.2. In the next step we derive the ex ante probability of the compromise, and the ex ante probability of any agent's most preferred alternative. Finally, we prove that the ex ante probabilities add up to more than one, which constitutes a contradiction.

Proof of Proposition 3.1: Suppose f is a first best decision rule. Then by definition for every agent i:

$$q_i(t_i) = 1 - G(1 - t_i).$$

We are going to use Lemma 3.3 to determine the conditional expected probability $p_i(t_i)$ that agent i's most preferred alternative is implemented. For this purpose, we first investigate agent i's conditional expected utility. By Lemma 3.3:

$$U_i(t_i) = U_i(1) - \int_{t_i}^{1} q_i(s_i)ds_i.$$

Now by the definition of first best $U_i(1) = 1$. Substituting this, and the formula for $q_i(t_i)$, we find:

$$U_i(t_i) = 1 - \int_{t_i}^{1} (1 - G(1 - s_i))\, ds_i.$$

Now recall the definition of conditional expected utility:

$$U_i(t_i) = p_i(t_i) + q_i(t_i)t_i.$$

We substitute the expressions for $q_i(t_i)$ and $U_i(t_i)$ that we have found above, and solve for $p_i(t_i)$:

$$p_i(t_i) = 1 - \int_{t_i}^{1} (1 - G(1 - s_i))\, ds_i - (1 - G(1 - t_i))t_i.$$

We simplify the formula on the right hand side:

$$p_i(t_i) = \int_{t_i}^{1} G(1 - s_i)ds_i + G(1 - t_i)t_i.$$

Finally we substitute the integration variable in the integral and obtain:

$$p_i(t_i) = \int_{0}^{1-t_i} G(s_i)ds_i + G(1 - t_i)t_i.$$

The next step is to calculate the ex-ante probability of the compromise, and i's most preferred alternative, calculated before i learns his type. These probabilities can be obtained by computing the expected values of the probabilities calculated above.

Thus, the ex-ante probability of the compromise is:

$$\begin{aligned} Q &\equiv \int_0^1 q_i(t_i)g(t_i)dt_i \\ &= \int_0^1 (1 - G(1-t_i))\, g(t_i)dt_i \\ &= 1 - \int_0^1 G(1-t_i)g(t_i)dt_i. \end{aligned}$$

The ex-ante probability of i's most preferred alternative is:

$$\begin{aligned} P &\equiv \int_0^1 \left(\int_0^{1-t_i} G(s_i)ds_i + G(1-t_i)t_i \right) g(t_i)dt_i \\ &= \int_0^1 \int_0^{1-t_i} G(s_i)g(t_i)ds_i dt_i + \int_0^1 G(1-t_i)g(t_i)t_i dt_i \end{aligned}$$

Changing the order of integration in the first expression, and integrating by parts in the second expression, we get:

$$\begin{aligned} P &= \int_0^1 \int_0^{1-s_i} G(s_i)g(t_i)dt_i ds_i + \int_0^1 G(1-t_i)g(t_i)t_i dt_i \\ &= \int_0^1 G(s_i) \int_0^{1-s_i} g(t_i)dt_i ds_i + \Big[G(1-t_i)G(t_i)t_i\Big]_0^1 \\ &\quad - \int_0^1 G(t_i)\left(G(1-t_i) - g(1-t_i)t_i\right) dt_i \\ &= \int_0^1 G(s_i)G(1-s_i)ds_i - \int_0^1 G(t_1)G(1-t_i)dt_i \\ &\quad + \int_0^1 G(t_i)g(1-t_i)t_i dt_i \\ &= \int_0^1 G(t_i)g(1-t_i)t_i dt_i \end{aligned}$$

The probability P calculated above is not only the ex-ante probability with which agent i's most preferred alternative is chosen, it is also the ex-ante probability with which the other agent's most preferred alternative is chosen. Thus, we have now determined for all three alternatives the ex-ante probability with which they are chosen. Clearly, these ex-ante alternatives have to add up to one. However, we shall now obtain a contradiction by showing:

$$Q + 2P > 1 \Leftrightarrow$$

$$1-\int_0^1 G(1-t_i)g(t_i)dt_i+2\int_0^1 (G(t_i)g(1-t_i)t_i)\,dt_i>1 \Leftrightarrow$$

$$\int_0^1 G(t_i)g(1-t_i)(2t_i-1)dt_i>0$$

To prove that this inequality is true, we denote the integral on the left hand side by I. Integration by parts yields:

$$\begin{aligned} I &= -\Big[G(t_i)G(1-t_i)(2t_i-1)\Big]_0^1 \\ &\quad +\int_0^1 G(1-t_i)\left(2G(t_i)+g(t_i)(2t_i-1)\right)dt_i \\ &= \int_0^1 2G(1-t_i)G(t_i)dt_i+\int_0^1 G(1-t_i)g(t_i)(2t_i-1)dt_i \\ &= \int_0^1 2G(1-t_i)G(t_i)dt_i - I \end{aligned}$$

Solving for I, we obtain:

$$I=\int_0^1 G(1-t_i)G(t_i)dt_i$$

It is obvious that this is strictly positive.

Q.E.D.

3.9 Discussion

Our main result is the impossibility derived in the previous section. We have shown that two individuals with opposing preferences who do not know the others' von Neumann Morgenstern preferences cannot compromise efficiently. Second best mechanisms can be determined as solutions to linear programming problems, but appear to be sensitive to the exogenous parameters.

Our study could be extended to a study of incentive compatible voting schemes. The voting problem differs from the compromising problem in that typically more than two people are involved. Moreover, not only the von Neumann Morgenstern utilities, but also the ordinal ranking of the alternatives by the other agents may be unknown.

We conjecture that the impossibility result that we found extends to the voting context, and that also in the voting context second best rules will be highly sensitive to the exogenous parameters of the model.

In practice rules for compromising, and voting rules, cannot be targeted for particular distributions of individual characteristics. They have to function well across a large variety of environments. To understand which rules for compromising, and which voting rules, are suitable in this context, one needs to develop a theory of "robust mechanism design". This appears to be the most important open problem in this area.[10]

[10] For some progress in this direction see, for example, Bergemann and Morris (2003)).

Chapter 4

Inefficient Procurement

4.1 Introduction

Many government tenders and procurement procedures not only require bidders to quote the price at which they are willing to supply a particular good or service, but also details of its specification. For example, the Strategic Rail Authority's (SRA) procedure for awarding rail franchises asks train operators to submit multi-dimensional bids, detailing proposals for infrastructure enhancements as well as the subsidy level required to put them in place. In their 2001 report, consultancy firm PriceWaterhouseCoopers raises concerns that the SRA's procedure may select a bidder whose enhancement proposals are "bad value for money" - the reason being that bidders must commit to a subsidy level at a time when the cost of implementing their proposed service specification is uncertain.[1] Do procedures that simultaneously determines price and specification always result in an efficient specification choice? If not, is there a justification for using such a procedure? These are the questions we address in this chapter.

There are various reasons why a procurer, instead of fully specifying the required good at the outset, may wish to use a two-dimensional procurement procedure. One reason is that the procurer may not know the possible specifications of the good, and

[1] *"Upgrading the Rail Network: Focusing on Delivery"*, PriceWaterhouse Coopers, August 2001 (see www.pwcglobal.com/rail).

therefore uses the procurement process to elicit this information. A second reason is that the buyer, while aware of all possible specifications, lacks the expertise to assess how much it costs to produce them. This is the view adopted in the present chapter, where we assume that suppliers have superior information, but also do not fully know the cost of the available specifications.

We develop this theme in a simple stylized model in which a buyer must purchase a good that exists in two specifications, both of which are considered equally good. There are two suppliers from whom the good can be sourced, and each of them can produce either specification. The suppliers' production cost for each specification are the same, but neither supplier actually knows them. Each of them has some private information about production cost in the form of two independent cost signals, namely one per specification. The production cost of a specification is simply the sum of the suppliers' signals about that specification.

Our setup captures the essence of the opening example of rail franchising, where it seems reasonable to assume that train operating companies are in a better position than the government to assess the cost of the various train technologies. However, the cost of implementing each train technology is likely to feature components that are common to all train operators: Tracks and signals in the UK are owned by NetworkRail, and hence, the implementation of any train technology depends on NetworkRail's ability to deliver the necessary track and signal work.

In this chapter, we first study a particular procurement procedure that essentially requires bidders to choose specification and price of the good in question. The two-dimensional bids are then evaluated according to the net benefit they generate for the buyer, with the winning bid being the one that yields the highest net benefit. As the buyer is indifferent between the two specifications, the suppliers' price quotes alone determine which supplier, and hence, which specification is chosen.

We show that this procedure may result in the choice of the specification whose production cost are higher. The reason is that, in equilibrium, each supplier chooses the specification for which he has observed the lower cost signal. Otherwise, expected production cost of a supplier with identical cost signals, conditional on winning, would be different for the two specifications.

We compare the inefficient procurement procedure to one that always selects the ex post efficient (i.e. cheapest) specification. Such a procedure requires the suppliers to report their cost signals, and then randomly chooses a supplier to produce the efficient specification. We show that under this procedure, the buyer has to make large lump-sum payments to induce the suppliers to truthfully disclose their cost information. The intuition is that each supplier has a chance of being selected to produce the buyer's chosen specification. Provided the payment he receives from the buyer is independent of his information, he has an incentive to ensure that the chosen specification is efficient and thus has lower production cost than the other.

A comparison of the (expected) payments that the buyer has to make in exchange for the good under the two procurement procedures leads us to conjecture that an optimal procurement mechanism will not generally be efficient.

The present work is related to the paper by Che (1993) who studies procurement auctions in a private value setting in which bidders submit two-dimensional bids specifying both quality (or specification) and price of the good. In contrast to our setting, he allows for a continuum of product specifications. Che studies scoring rules that assign a one-dimensional score to every (two-dimensional) bid. He considers a variety of procedures for choosing a supplier and specification on the basis of the awarded scores. One such procedure is the first score auction, in which the bidder whose bid has the highest score wins the auction, and supplies the good in the specification and at the price stated in his bid. In our setting, in which the two available specifications are equally good for the buyer, the first score auction corresponds to our procurement procedure that selects specification and supplier on the basis of price alone. In Che's model, the first score auction reduces to the standard first price auction with private values, and therefore always awards the contract to the most efficient supplier. The distinguishing feature of our model is that suppliers' production cost are "common value" and that their signal spaces are two-dimensional. We show that in our setup the first score auction is not efficient. This also contrasts with the efficiency result for the first price auction in the standard symmetric common value setup (see Maskin (1992)).

The present work is also related to the paper by Jehiel and Moldovanu (2001) who study the question whether efficiency can be attained in general social choice settings where the dimension of the agents' signal spaces is the same as the number of possible outcomes. In such settings, a one-dimensional payment (per agent and alternative) should in principle suffice to elicit the entire information. However, if the agents' payoffs/valuations are interdependent, efficient and incentive compatible direct revelation mechanisms exist only if a very restrictive symmetry condition holds. This condition is trivially satisfied in our model. To see why this is the case, note that due to the "pure common value" nature of production cost, it is irrelevant which supplier produces the efficient specification. Thus, the buyer can select a supplier at random, whose expected production cost of each specification are therefore a fixed proportion of the actual production cost.[2] This allows the buyer to align individual incentives with the goal of achieving ex post efficiency.

The remainder of this chapter is structured as follows: Section 4.2 contains the description of the model. Section 4.3 is studies the procurement procedure under which suppliers' two-dimensional bids contain specification and price of the good. Section 4.4 details an efficient procurement procedure, and Section 4.5 concludes with a brief discussion of the findings. Section 4.6 (Appendix 1) contains a characterization of suppliers' expected payments under any incentive compatible procurement mechanism. Section 4.7 (Appendix 2) contains some proofs.

4.2 Model

There is a buyer who must buy one unit of a particular good. There are two suppliers from whom this good can be sourced, and it can be produced either in specification A or in specification B. We consider the limiting case in which the buyer derives the

[2]The expectation is taken with respect to the probability of being chosen by the buyer.

same benefit b from both specifications.

The true cost of producing specifications A and B, c_A and c_B respectively, are unknown at the time at which suppliers compete for the buyer's custom. However, each supplier has some private information about the true value of c_A and c_B. Supplier i's private information consists of two signals: s_A^i and s_B^i ($i = 1, 2$). The signal s_A^i is i's private information about the cost of specification A, and the signal s_B^i is i's private information about the cost of specification B. Each signal is a random variable with domain $[0, 1]$. We assume that all signals are identically and independently distributed according to some distribution G with density g.

The production cost of each specification are the same across suppliers.[3] Furthermore, we assume that c_A can be expressed as the sum of the suppliers' signals about specification A, and that c_B can be expressed as the sum of the suppliers' signals about specification B:

$$c_A(s_A^1, s_A^2) = s_A^1 + s_A^2 \text{ and } c_B(s_B^1, s_B^2) = s_B^1 + s_B^2.$$

The buyer can either purchase the good from supplier 1 or from supplier 2, and in each case, either in specification A or in specification B. The buyer's von Neumann Morgenstern utility function is $b - p$, where p is the price she pays for the good. Supplier i's von Neumann Morgenstern utility function is $p - c_k$ if the good is purchased from supplier i in specification k ($k \in \{A, B\}$) at price p, and zero if the good is purchased from supplier j ($i, j = 1, 2$, $j \neq i$).

The buyer proposes the suppliers a procurement procedure with the aim of simultaneously determining the specification and price of the good, as well as the identity of the supplier. Based on their private information about production cost, the suppliers can either participate in the buyer's procedure or opt out. If a supplier opts out, he receives his outside option, which is assumed to yield him a payoff of zero. In the following two sections, we analyze two specific procurement procedures.

[3]Thus, this is a "pure common value" model.

4.3 Minimum Price Mechanism

4.3.1 Setup

Suppose the buyer asks each supplier i ($i = 1, 2$) to submit a pair of prices: a price p^i_A to be paid if the buyer sources the good from supplier i in specification A, and a price p^i_B to be paid if the buyer sources the good from supplier i in specification B. The buyer commits to sourcing the good at the lowest of all four prices. That is, the buyer purchases the good from the supplier whose minimum price is lowest. If the minimum price of supplier 1 is the same as the minimum price of supplier 2, then each supplier has equal probability of being selected by the buyer. We call this procurement procedure *minimum price mechanism* (MPM).

Formally, the MPM, in conjunction with the suppliers' signal spaces $S = [0,1]^2$, the probability distribution G, and the cost functions c_A and c_B, constitutes a Bayesian game of incomplete information. The structure of the game is as follows:

- **Stage 1**: Supplier 1 privately observes his signal-vector $\mathbf{s}^1 = (s^1_A, s^1_B) \in S$, and simultaneously supplier 2 privately observes his signal-vector $\mathbf{s}^2 = (s^2_A, s^2_B) \in S$.

- **Stage 2**: Based on $\mathbf{s}^1$ supplier 1 submits a pair of prices $(p^1_A, p^1_B) \in \mathbb{R}^2_+$. Simultaneously supplier 2 submits a bid $(p^2_A, p^2_B) \in \mathbb{R}^2_+$ based on his signal $\mathbf{s}^2$.

- **Stage 3**: The buyer sources the good from supplier i^*, where $i^* \in \arg\min_{i\in\{1,2\}}(\min_{k\in\{A,B\}} p^i_k)$. Supplier i^* must provide the good in specification k^*, where $k^* \in \arg\min_{k\in\{A,B\}} p^{i^*}_k$. Supplier i^* receives a payment equal to his minimum price $p^{i^*}_{k^*}$ from the buyer and incurs production cost c_{k^*}. Supplier j ($j \in \{1,2\}$, $j \neq i^*$) receives no payment and incurs no cost. The buyer's payoff is $b - p^{i^*}_{k^*}$.

A strategy for supplier i consists of two pricing functions, one per specification. Formally, $p^i_k : S \to \mathbb{R}_+$, $\mathbf{s}^i \mapsto p^i_k(\mathbf{s}^i)$ for every $k \in \{A, B\}$. In order to facilitate the characterization of Bayes Nash Equilibria (BNE) of this game, we introduce the following alternative version of the MPM:

Suppose the buyer asks each supplier to submit a two-dimensional bid consisting of a specification and a price. Let $k^i \in \{A, B\}$ denote the specification chosen by supplier i, and let $p^i_{k^i}$ denote the price at which i is willing to provide specification k^i. The buyer commits to sourcing the good at the lowest price, irrespective of the proposed specification(s): Given prices $p^1_{k^1}$ and $p^2_{k^2}$, the buyer chooses the good from supplier i^*, where

$$i^* \in \arg\min_{i \in \{1,2\}} p^i_{k^i}.$$

Supplier i^* is paid his price $p^{i^*}_{k^{i^*}}$ in exchange for the good in specification k^{i^*}. In producing k^{i^*}, he incurs cost $c_{k^{i^*}}$. Supplier j ($j \in \{1, 2\}$, $j \neq i^*$) receives no payment and incurs no cost. The buyer's payoff is $b - p^{i^*}_{k^{i^*}}$.

A strategy for supplier i in this alternative version of the MPM consists of *three* functions - the specification choice rule k^i and the pricing functions p^i_A and p^i_B:

$$k^i : S \to \{A, B\},\ \mathbf{s}^i \mapsto k^i(\mathbf{s}^i),$$

$$p^i_l : S \to \mathbb{R}_+,\ \mathbf{s}^i \mapsto p^i_l(\mathbf{s}^i) \text{ for every } l \in \{A, B\}.$$

The interpretation is as follows: If $k^i(\mathbf{s}^i) = l$ then supplier i offers to produce the good in specification l at price $p^i_l(\mathbf{s}^i)$.

In the next section, we characterize a symmetric BNE of the alternative MPM. The following Lemma states that this BNE will also be a BNE of the original MPM. This is intuitive, as a supplier's decision *not* to offer a particular specification under the alternative MPM can be replicated in the original MPM by setting a price that is strictly above the one for the other specification.

Lemma 4.1 *Every BNE of the alternative MPM constitutes a BNE of the original MPM.*

The proof is relegated to Section 4.7 (Appendix 2).

4.3.2 Equilibrium Characterization

In this section, we characterize a BNE of the alternative MPM introduced in the previous section. We make the following assumptions about the suppliers' equilibrium strategies:

Assumption 4.1 (Symmetry w.r.t. suppliers) *Both suppliers use the same strategy. Thus, if the two suppliers have identical cost signals, they choose the same specification and quote the same price for it. Formally, if* $s_A^1 = s_A^2 = x$ *and* $s_B^1 = s_B^2 = y$ *it holds that* $k^1(x,y) = k^2(x,y) \equiv k(x,y)$ *and* $p_l^1(x,y) = p_l^2(x,y) \equiv p_l(x,y)$ *for every* $l \in \{A, B\}$ *and all* $(x,y) \in S$.

Assumption 4.2 (Specification choice rule) *The specification choice rule* k *takes the form*

$$k(s_A^i, s_B^i) = \begin{cases} A & \text{if } s_B^i > I(s_A^i) \\ B & \text{if } s_B^i \leq I(s_A^i) \end{cases} \quad \textit{for all } (s_A^i, s_B^i) \in S, \tag{4.1}$$

where $I : [0,1] \rightarrow [0,1]$, $s_A^i \mapsto I(s_A^i)$ *is a strictly increasing function.*

Assumption 4.3 (Separability) *The price quoted for any specification depends only on the signal realization pertaining to that specification. That is, for all* $(s_A^i, s_B^i) \in S$ *it holds that* $p_A(s_A^i, s_B^i) \equiv p_A(s_A^i)$ *and* $p_B(s_A^i, s_B^i) \equiv p_B(s_B^i)$.

Assumption 4.4 (Symmetry w.r.t. specifications) *The pricing functions for the two specifications are identical:* $p_A(s) = p_B(s) \equiv p(s)$ *for all* $s \in [0,1]$.

Assumption 4.5 (Monotonicity) *The pricing function* p *is increasing and differentiable on* $[0,1)$.

We can show that in a Bayes Nash Equilibrium (k,p) that satisfies Assumptions 4.1-4.5, each supplier chooses his specification on the basis of his *minimum signal.* I.e. the specification choice rule in (4.1) is such that $I(s_A^i) = s_A^i$ for all $s_A^i \in [0,1]$. For this purpose, suppose there exists an interval $Z \subset [0,1]$ such that $z > I(z)$ for all $z \in Z$.

Consider supplier 1 and suppose he has identical signal realizations $s_A^1 = s_B^1 = z$. He will therefore offer specification A at price $p(z)$. Figure 4.1 (a) provides an illustration.

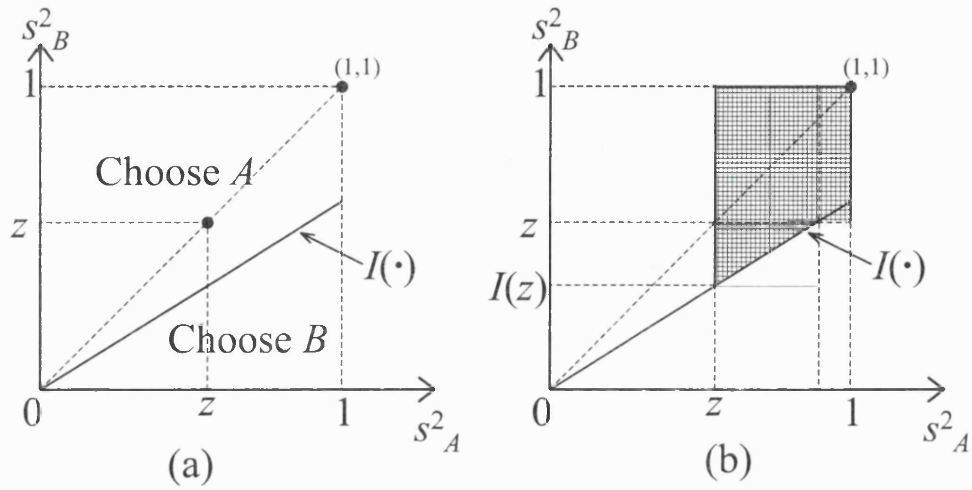

Figure 4.1: Specification Choice

If supplier 1 wins the contract, he has to provide the good in specification A in return for a payment of $p(z)$, while his production cost are $z + s_A^2$. The shaded area in Figure 4.1 (b) shows all signal realizations (s_A^2, s_B^2) of supplier 2 for which supplier 1 wins the contract. This is the case in the following events:

1. Supplier 2 offers the good in specification A at a price $p(s_A^2) > p(z)$. This implies that $s_B^2 > I(s_A^2)$ and $s_A^2 > z$.

2. Supplier 2 offers the good in specification B at a price $p(s_B^2) > p(z)$. This implies that $s_B^2 \leq I(s_A^2)$ and $s_B^2 > z$.

It is easy to show that, conditional on winning the contract at price $p(z)$, the expected production cost of specification A are strictly higher than those of specification B. Thus, supplier 1 would prefer to offer the good in specification B rather than A, which constitutes a contradiction. To see this, note that, conditional on supplier 1 winning the contract, the expected value of supplier 2's signal about specification A is

$$(1 - G(z))^2 E[s_A^2 \,|s_A^2 > z] + \int_z^{I^{-1}(z)} \left(G(z) - G(I(s_A^2))\right) s_A^2 g(s_A^2) ds_A^2. \tag{4.2}$$

The expected value of supplier 2's signal about specification B, conditional on supplier 1 winning the contract, is

$$(1-G(z))^2 E[s_B^2 \,|s_B^2 > z] + \int_z^{I^{-1}(z)} \left(G(z) - G(I(s_A^2))\right) E[s_B^2 \,|I(s_A^2) \le s_B^2 < z] g(s_A^2) ds_A^2. \tag{4.3}$$

As s_A^2 and s_B^2 are i.i.d. random variables, the difference in expected value between (4.2) and (4.3) is

$$\int_z^{I^{-1}(z)} \left(G(z) - G(I(s_A^2))\right) [s_A^2 - E[s_B^2 \,|I(s_A^2) \le s_B^2 < z]] g(s_A^2) ds_A^2.$$

As $E[s_B^2 \,|I(s_A^2) \le s_B^2 < z] < z$ and $s_A^2 \in [z, I^{-1}(z)]$, it follows immediately that this difference is strictly positive. Thus, conditional on winning, specification B has lower production cost than specification A.[4] The following result completes the equilibrium characterization:

Proposition 4.1 *The unique Bayes Nash Equilibrium (k, p) of the alternative MPM that satisfies Assumptions 4.1 - 4.5 is given by the specification choice rule*

$$k(s_A^i, s_B^i) = \begin{cases} A & \text{if } s_B^i > s_A^i \\ B & \text{if } s_B^i \le s_A^i \end{cases} \quad \textit{for all } (s_A^i, s_B^i) \in S,$$

and the pricing function

$$p(z^i) = \frac{1}{(1-G(z^i))^2} \left(\int_{z^i}^1 \left(2t + \int_t^1 (t+s) \frac{g(s)}{1-G(t)} ds \right) g(t)(1-G(t)) dt \right),$$

where $z^i = \min\{s_A^i, s_B^i\}$.

The remainder of this section is devoted to the derivation of the equilibrium pricing function in Proposition 4.1.

[4] A similar argument shows that it cannot be the case that $z < I(z)$.

Proof of Proposition 4.1:

Supplier 1's maximization problem: Consider supplier 1 and suppose that $s_A^1 < s_B^1$. Suppose supplier 2 uses the equilibrium strategy (k, p) in Proposition 4.1 and that supplier 1, instead of submitting the price $p(s_A^1)$, quotes some other price for his chosen specification.

Notice that it does not pay for supplier 1 to quote a price p_1 less than $p(0)$ since in that case, supplier 1 would be chosen for sure, but could do better by increasing his price slightly so that he is still chosen for sure but has a higher payoff. Second, any price p_1 above $p(1)$ means that supplier 2 is chosen for sure, leaving supplier 1 with a profit of zero. The same outcome can be achieved by setting $p_1 = p(1)$. Thus, we only need to consider prices p_1 in $[p(0), p(1)]$.

Quoting a price $p_1 \in [p(0), p(1)]$ is equivalent to choosing some cost signal $s \in [0, 1]$ (where s need not be equal to supplier 1's minimum signal s_A^1) and quoting the corresponding equilibrium price $p(s)$. In the following, we assume that supplier 1 with minimum signal s_A^1 quotes price $p(s)$ for specification A. Conditional on winning, he receives a payment equal to $p(s)$. He wins if he charges a lower price than supplier 2, which is the case in all events in which $\min\{s_A^2, s_B^2\} > s$. Supplier 1's expected profit therefore is

$$(1 - G(s)) \left((1 - G(s))p(s) - \int_s^1 (s_A^1 + s_A^2) g(s_A^2) ds_A^2 \right). \tag{4.4}$$

Supplier 1 solves the following problem:

$$\max_s (1 - G(s)) \left((1 - G(s))p(s) - \int_s^1 (s_A^1 + s_A^2) g(s_A^2) ds_A^2 \right).$$

The first order condition is

$$- g(s) \left((1 - G(s))p(s) - \int_s^1 (s_A^1 + s_A^2) g(s_A^2) ds_A^2 \right) \\ + (1 - G(s)) \left(-g(s)p(s) + (1 - G(s))p'(s) + (s_A^1 + s) g(s) \right) \stackrel{!}{=} 0.$$

In equilibrium, the optimal value of s must be equal to s_A^1, and so setting $s = s_A^1$ in the first-order condition, we obtain the following linear differential equation of first order

$$p'(s_A^1) + p(s_A^1)\left(\frac{-2g(s_A^1)}{1-G(s_A^1)}\right) = -\left(2s_A^1 + \int_{s_A^1}^1 (s_A^1 + s_A^2)\frac{g(s_A^2)}{1-G(s_A^1)}ds_A^2\right)\frac{g(s_A^1)}{1-G(s_A^1)}. \tag{4.5}$$

General solution to the differential equation in (4.5): In order to obtain an explicit expression for the pricing function p we need to solve the differential equation in (4.5). The first step is to multiply the equation through by $(1-G(s_A^1))^2$. This yields

$$\begin{aligned}(1-G(s_A^1))^2 p'(s_A^1) &+ p(s_A^1)\left(-2g(s_A^1)(1-G(s_A^1))\right) \\ &= -\left(2s_A^1 + \int_{s_A^1}^1 (s_A^1 + s_A^2)\frac{g(s_A^2)}{1-G(s_A^1)}ds_A^2\right) g(s_A^1)(1-G(s_A^1)).\end{aligned}$$

Note that the left-hand side of the latter equation is the derivative of $(1-G(s_A^1))^2 p(s_A^1)$. Integrating both sides of the latter equation from s_A^1 to 1 yields the unique solution

$$-(1-G(s_A^1))^2 p(s_A^1) = -\int_{s_A^1}^1 \left(2t + \int_t^1 (t + s_A^2)\frac{g(s_A^2)}{1-G(t)}ds_A^2\right) g(t)(1-G(t))dt$$

or

$$p(s_A^1) = \frac{1}{(1-G(s_A^1))^2}\left(\int_{s_A^1}^1 \left(2t + \int_t^1 (t + s_A^2)\frac{g(s_A^2)}{1-G(t)}ds_A^2\right) g(t)(1-G(t))dt\right). \tag{4.6}$$

Boundary condition: Consider supplier 1 and suppose that he has a minimum signal of 1 (i.e. $s_A^1 = s_B^1 = 1$). Assume that supplier 2 uses the equilibrium strategy (k, p). Conditional on winning (which happens in the event that supplier 2 has the same signal realizations as supplier 1) the (expected) production cost of supplier 1 are $c_A = c_B = 2$.

Suppose now that the equilibrium pricing function is such that $p(1) > 2$. Figure 4.2 (a) provides an illustration. In the event that supplier 1 wins the contract he has a strictly positive payoff. However, the probability that supplier 1 is chosen when he quotes the price $p(1)$ (which is $\Pr\{s_A^2 \geq 1 \wedge s_B^2 \geq 1\}$) is zero. Thus, his expected payoff

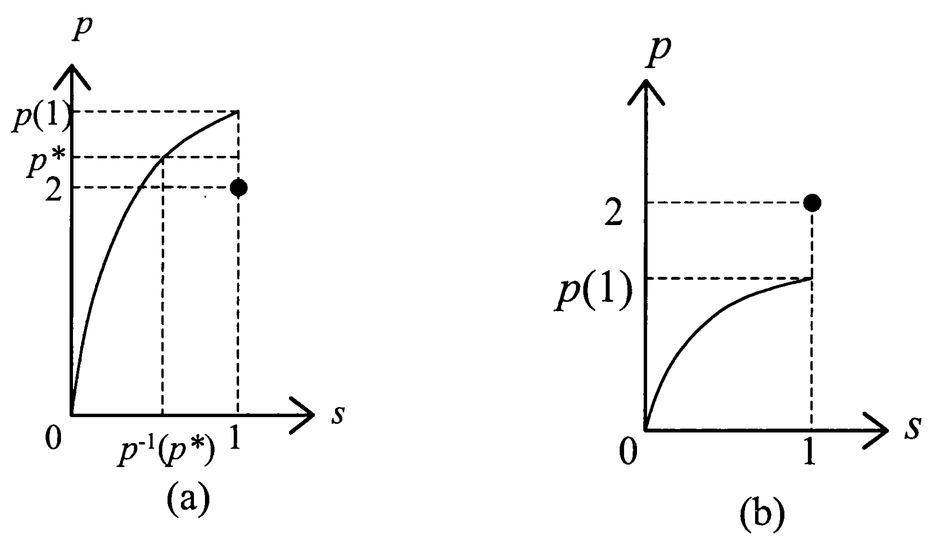

Figure 4.2: Boundary Condition

is zero. Now consider the following deviation: supplier 1 offers specification A, say, and charges the price p^*, where $2 < p^* < p(1)$. This results in an expected payoff of

$$\int_{p^{-1}(p^*)}^{1}\int_{p^{-1}(p^*)}^{1}(p^* - 1 - s_A^2)g(s_A^2)g(s_B^2)ds_A^2 ds_B^2,$$

or equivalently

$$(1 - G(p^{-1}(p^*)))^2(p^* - 1 - E[s_A^2\,|s_A^2 > p^{-1}(p^*)]).$$

This expected payoff is strictly positive as $p^* > 2 > 1 + E[s_A^2\,|s_A^2 > p^{-1}(p^*)]$. This establishes a contradiction.

Now suppose that the pricing function is such that $p(1) < 2$. Figure 4.2 (b) provides an illustration. Suppose that supplier 1's minimum signal is $1 - \varepsilon$, where $\varepsilon > 0$ but small. In the event that supplier 1 wins the contract (which occurs with probability $(1 - G(1 - \varepsilon))^2$) he has a strictly negative payoff. Note that it is profitable for supplier 1 to deviate to $p^* = 2$. In the limit as $\varepsilon \to 0$ supplier 1 charges the price $p^* = 2$ and has a profit of zero. Thus, the boundary condition associated with (4.5) is

$$p(1) = 2.$$

Note that the differential equation in (4.5) has a unique solution (namely the function p given in (4.6)). We can easily verify that this function p satisfies the boundary condition derived above. As $(1-G(s_A^1))^2 = 0$ at $s_A^1 = 1$ we know that p is not defined at this point. We therefore compute $\lim_{s_A^1 \to 1} p(s_A^1)$. To compute the limit, we use L'Hôpital's Rule. For this purpose note that the numerator of the pricing function $p(\cdot)$ in (4.6) can be equivalently written as

$$\int_{s_A^1}^1 3tg(t)(1-G(t))dt + \int_{s_A^1}^1 E[s_A^2\,|s_A^2 \geq t]g(t)(1-G(t))dt.$$

It's derivative w.r.t. s_A^1 (via Leibniz' Rule) is

$$-3s_A^1 g(s_A^1)(1-G(s_A^1)) - E[s_A^2\,|s_A^2 \geq s_A^1]g(s_A^1)(1-G(s_A^1)).$$

Now take the derivative of the denominator of the pricing function $p(\cdot)$ in (4.6). This yields $-2g(s_A^1)(1-G(s_A^1))$. Dividing the derivatives of numerator and denominator, we obtain

$$\frac{3}{2}s_A^1 + \frac{1}{2}E[s_A^2\,|s_A^2 \geq s_A^1].$$

Thus, $\lim_{s_A^1 \to 1}\left(\frac{3}{2}s_A^1 + \frac{1}{2}E[s_A^2\,|s_A^2 \geq s_A^1]\right) = 2$, as $E[s_A^2\,|s_A^2 \geq 1] = 1$.

Proof that if supplier 2 uses $p(\cdot)$ then it is optimal for supplier 1 to do so: We now show that it does not pay for supplier 1 (whose minimum signal is s_A^1) to "pretend" that his minimum signal is $s < s_A^1$: The expected payoff if supplier 1 submits the price $p(s)$ when his true minimum signal is s_A^1 is

$$\Pi_1(s, s_A^1) = (1-G(s))\left((1-G(s))p(s) - \int_s^1 (s_A^1 + s_A^2)g(s_A^2)ds_A^2\right).$$

Hence,

$$\begin{aligned}\frac{\partial \Pi_1(s, s_A^1)}{ds} &= -2g(s)(1-G(s))p(s) + (1-G(s))^2 p'(s)\\ &\qquad + g(s)(1-G(s))(2s_A^1 + s) + g(s)\int_s^1 s_A^2 g(s_A^2)ds_A^2.\end{aligned}$$

Note that $2s_A^1 + s > 2s + s$ as we assumed that $s_A^1 > s$. We can easily compute $\Pi_1(z, s)$ and $\frac{\partial \Pi_1(z,s)}{dz}$. As $p(\cdot)$ maximizes the expected utility of a supplier whose minimum signal is s, it holds that

$$\left.\frac{\partial \Pi_1(z,s)}{dz}\right|_{z=s} = -2g(s)(1-G(s))p(s) + (1-G(s))^2 p'(s)$$
$$+ g(s)(1-G(s))(3s) + g(s)\int_s^1 s_A^2 g(s_A^2) ds_A^2$$
$$= 0.$$

As $2s_A^1 + s > 2s + s = 3s$ we have

$$\frac{\partial \Pi_1(s, s_A^1)}{ds} > \left.\frac{\partial \Pi_1(z,s)}{dz}\right|_{z=s} = 0$$

or $\frac{\partial \Pi_1(s,s_A^1)}{ds} > 0$. Thus, supplier 1 can raise his expected payoff by raising his "report" s. An analogous argument shows that if $s > s_A^1$ it holds that $\frac{\partial \Pi_1(s,s_A^1)}{ds} < 0$. Therefore, supplier 1's expected payoff $\Pi_1(s, s_A^1)$ is maximized by choosing $s = s_A^1$.

Q.E.D.

Note that under the equilibrium pricing function in Proposition 4.1, each supplier's ex ante expected payment under the alternative MPM is given by

$$2\int_0^1 p(s)(1-G(s))^3 g(s) ds.$$

As an example, consider the case in which all signals are distributed uniformly on [0, 1]. We obtain $p(z^i) = \frac{7}{6}z^i + \frac{5}{6}$, and hence the ex ante payment is

$$2\int_0^1 (\frac{7}{6}s + \frac{5}{6})(1-s)^3 ds = \frac{8}{15}.$$

Thus, the expected expenditure that the buyer has to make in order to procure the good using the MPM is $16/15 \approx 1.067$.

4.3.3 Inefficiency of Minimum Price Mechanism

The outcome of the procurement process is (ex post) *efficient* if and only if the buyer sources the good in the specification for which production cost are lowest. Note that efficiency of the outcome of the procurement process pertains only to the choice of specification, *not* to the choice of supplier: Due to the symmetry of the setup, the buyer does not care from which supplier he sources the good, as long as it is provided in the efficient specification.

Recall that under the minimum price mechanism the winner is the supplier who has the lowest minimum signal, *irrespective* of the specification to which this signal pertains. It is not surprising that there are signal constellations for which the outcome is inefficient. Suppose that $s_B^2 < s_A^2$. Then, the minimum price mechanism is *inefficient* if:

Case 1: Specification A is implemented although specification B is efficient. This is the case for all pairs (s_A^1, s_B^1) such that

$$s_A^1 < s_B^2 \text{ and } s_A^1 \leq s_B^1 < s_A^1 + (s_A^2 - s_B^2),$$

resulting in supplier 1 being chosen.

Case 2: Specification B is implemented although specification A is efficient. This is the case for all pairs (s_A^1, s_B^1) such that

$$s_B^2 \leq s_A^1 \leq 1 - (s_A^2 - s_B^2) \text{ and } s_B^1 > s_A^1 + (s_A^2 - s_B^2),$$

resulting in supplier 2 being chosen.

The shaded areas in Figure 4.3 represent all pairs (s_A^1, s_B^1) for which the minimum price mechanism selects the good's specification efficiently. The unshaded regions highlight the inefficiencies.[5] Figure 4.3 also shows that for all pairs (s_A^1, s_B^1) for which the minimum price mechanism results in an inefficient specification choice, it must hold that $s_B^1 > s_A^1$. This illustrates a more general point: A *necessary* condition for an inefficient outcome is that the specification associated with supplier 1's minimum signal is *not* the same as the specification associated with supplier 2's minimum signal. Consequently, if $s_A^i < s_B^i$ and the outcome is inefficient, then it must hold that $s_A^j > s_B^j$ $(i, j = 1, 2,\ i \neq j)$. It is easy to see that the minimum price mechanism is efficient whenever the suppliers' minimum signals pertain to the same specification: Suppose

[5]The figure is based on $s_A^2 = 0.8$ and $s_B^2 = 0.6$.

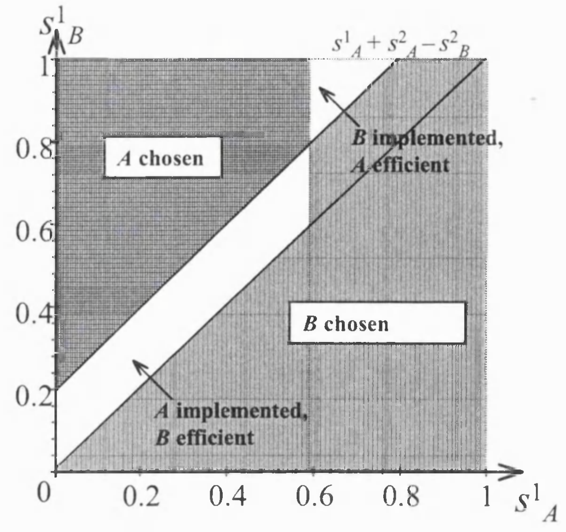

Figure 4.3: Inefficiency of MPM

that each suppliers' minimum signal pertains to specification A, so that the good is produced in specification A (by supplier i^*, where $i^* = \arg\min_{i\in\{A,B\}} s^i_A$). The cost of producing specification A is $c_A = s^1_A + s^2_A$. As $s^2_A < s^2_B$ it follows immediately that $c_A < s^1_A + s^2_B$. As we also have $s^1_A < s^1_B$, it then follows that $s^1_A + s^2_B < s^1_B + s^2_B = c_B$. Thus, the cost of specification A are lower than those of specification B.

This argument also illustrates why inefficiencies arise under the minimum price mechanism: In order to determine the efficient specification in situations where the minimum signals of the two suppliers pertain to different specifications, we must not only look at the suppliers' minimum signals, but also at the *difference* between the two signals of each supplier. This information, however, is ignored by the minimum price mechanism. In the next section, we discuss a mechanism that achieves efficiency.

4.4 Efficient Mechanism

Recall that specification A is efficient if and only if $s^1_A + s^2_A < s^1_B + s^2_B$. For given signal values s^1_A and s^1_B (where $s^1_A > s^1_B$), Figure 4.4 shows all signal pairs (s^2_A, s^2_B) for which specifications A and B must be chosen in order to guarantee efficiency.

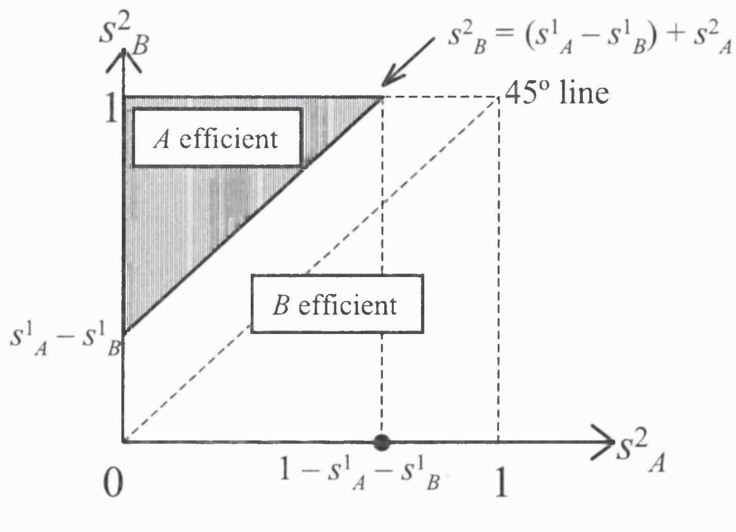

Figure 4.4: Efficient Specification Choice

As production cost are the same for both suppliers, it is irrelevant which supplier produces the efficient specification. Therefore, consider the following direct revelation mechanism:

Each supplier is asked to report his cost signals in return for a fixed payment. The buyer selects a supplier using a random device that selects supplier i ($i = 1, 2$) with probability λ_i ($0 \leq \lambda_i \leq 1$ and $\lambda_1 + \lambda_2 = 1$). The chosen supplier is then required to produce the good in the efficient specification. No supplier can be forced to participate in the mechanism. If a supplier decides not to participate, he receives his outside option payoff of zero.

Formally, if the suppliers' cost signals are s^1_A, s^1_B, s^2_A and s^2_B, let $\pi^{eff}_{Ai}(s^1_A, s^1_B, s^2_A, s^2_B)$ be the probability that supplier i ($i = 1, 2$) is selected to produce the good in specification A:

$$\pi^{eff}_{Ai}(s^1_A, s^1_B, s^2_A, s^2_B) = \begin{cases} \lambda_i & \text{if } s^1_A + s^2_A < s^1_B + s^2_B \\ 0 & \text{if } s^1_A + s^2_A > s^1_B + s^2_B \end{cases} \tag{4.7}$$

Similarly, define as $\pi^{eff}_{Bi}(s^1_A, s^1_B, s^2_A, s^2_B)$ the probability that supplier i is selected to produce the good in specification B:

$$\pi_{Bi}^{eff}(s_A^1, s_B^1, s_A^2, s_B^2) = \begin{cases} 0 & \text{if } s_A^1 + s_A^2 < s_B^1 + s_B^2 \\ \lambda_i & \text{if } s_A^1 + s_A^2 > s_B^1 + s_B^2 \end{cases} \tag{4.8}$$

Finally, let τ_i^{eff} be the fixed payment that the buyer pays supplier i for reporting his signals:

$$\tau_i^{eff} = \lambda_i \left(1 + 2\int_0^1 s(1 - G(s))g(s)ds\right).$$

It is easy to verify that τ_i^{eff} is equal to the (interim) expected production cost of supplier i when both his cost signals are 1.

The suppliers' cost signals are private information, and therefore the buyer must ensure that supplier 1, when invited to report his signal vector $\mathbf{s}^1$, has no incentive to lie:

$$\mathbf{s}^1 \in \arg\max_{\widehat{\mathbf{s}}^1 \in S}(\tau_1^{eff} - C_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1)).$$

where $C_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1)$ denotes the expected production cost of supplier 1 when his signal vector is $\mathbf{s}^1$ but he report it as $\widehat{\mathbf{s}}^1$:

$$C_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1) = \int_S\int (\pi_{A1}^{eff}(\widehat{\mathbf{s}}^1, \mathbf{s}^2)c_A(s_A^1, s_A^2) + \pi_{B1}^{eff}(\widehat{\mathbf{s}}^1, \mathbf{s}^2)c_B(s_B^1, s_B^2))g(s_A^2)g(s_B^2)ds_A^2 ds_B^2.$$

Similarly, supplier 2, when invited to report his signals $\mathbf{s}^2$, must not have any incentive to lie:

$$\mathbf{s}^2 \in \arg\max_{\widehat{\mathbf{s}}^2 \in S}(\tau_2^{eff} - C_2(\widehat{\mathbf{s}}^2, \mathbf{s}^2)).$$

Supplier 2's expected production cost $C_2(\widehat{\mathbf{s}}^2, \mathbf{s}^2)$ are defined analogously to those of supplier 1. It is straightforward to show the following:

Proposition 4.2 *The efficient mechanism* $\{\pi_{Ai}^{eff}, \pi_{Bi}^{eff}, \tau_i^{eff}\}_{i=1,2}$ *has a Bayes Nash Equilibrium in which each supplier participates and truthfully reveals his cost signals.*

This result follows immediately from the fact that if supplier i ($i = 1, 2$) truthfully reveals his signals, then supplier j ($j = 1, 2$, $j \neq i$) minimizes his expected production cost $C_j(\widehat{\mathbf{s}}^j, \mathbf{s}^j)$ by reporting his signals truthfully.[6] To see this, note that each supplier

[6]This is shown formally in Appendix 2.

has a positive chance of being chosen to produce whichever specification is efficient. Hence no supplier has an incentive to misrepresent his cost signals, as this may result in him having to produce the more costly of the two specifications. By the same logic, it holds that an efficient direct revelation mechanism is incentive compatible if and only if the transfer to each agent is a constant. The need to ensure participation then determines the level of the transfer. For example, if all signals are uniformly distributed on $[0, 1]$, we have $C_i(1, 1, 1, 1) = \frac{2}{3}\lambda_i$, so that the buyer has to pay a total of $\frac{4}{3}$ for the good.

4.5 Discussion

The main contribution of this chapter is the characterization of suppliers' equilibrium strategies in the Bayesian game induced by the Minimum Price Mechanism. We have shown that this mechanism is not efficient. With uniformly distributed signals, the buyer's expected payment for the good is roughly 1.067, which is strictly lower than expenditure under the efficient direct revelation mechanism. We can therefore conclude that the efficient mechanism given in (4.7) and (4.8) is not optimal (i.e. does not minimize the buyer's expected expenditure for the good). We also conjecture that efficient direct revelation mechanisms in which each suppliers' probability λ_i $(i = 1, 2)$ is a function of the reports (i.e. $\lambda_i(s_A^1, s_B^1, s_A^2, s_B^2)$) will not allow the buyer to achieve efficiency more "cheaply" than with constant λ_i. The reason is that incentive compatibility imposes stringent conditions on how λ_i can depend on supplier i's signals.

A characterization of expenditure minimizing incentive compatible mechanisms (even in the uniform case) is difficult due to the multi-dimensionality of the suppliers' signals. In Appendix 1 to this chapter, we derive the suppliers' equilibrium payoff functions associated with incentive compatible mechanisms. We thereby obtain a "revenue equivalence" result for our setup with two-dimensional private information, stating that the suppliers' payoff functions from any two incentive compatible procurement mechanisms with the same conditional expected allocation probabilities differ at most by an additive constant.[7] This result forms the basis for a characterization of optimal mechanisms, which is left for future work.

Finally, note that the existence of an efficient direct revelation mechanism in our setup hinges on the fact that there is (at least) one supplier who can produce whichever specification is efficient. In more general versions of our setup (i.e. those in which production cost for each specification are interdependent, but need not be the same across suppliers), an efficient and incentive compatible mechanism exists if and only if it is possible to separate the question of which specification is efficient from the question which of the two suppliers should produce it. The following example illustrates a situation where this is not possible:

Suppose the suppliers' production cost are as follows: $c_A^1(s_A^1, s_A^2) = s_A^1 + s_A^2$, $c_A^2(s_A^1, s_A^2) = 2\,(s_A^1 + s_A^2)$, $c_B^1(s_B^1, s_B^2) = 2\,(s_B^1 + s_B^2)$ and $c_B^2(s_B^1, s_B^2) = s_B^1 + s_B^2$. If $s_A^1 + s_A^2 < s_B^1 + s_B^2$ it is efficient to choose supplier 1 to produce specification A. Otherwise it is efficient to choose supplier 2 to produce specification B. From the point of view of achieving efficiency, we can recast this setup as one in which there are only two possible outcomes: either supplier 1 is chosen to produce A or supplier 2 is chosen to produce specification B.

The suppliers' induced cost functions are therefore $C_A^1(s_A^1, s_A^2) \equiv c_A^1(s_A^1, s_A^2) = s_A^1 + s_A^2$, $C_A^2(s_A^1, s_A^2) = 0$, $C_B^1(s_B^1, s_B^2) = 0$ and $C_B^2(s_B^1, s_B^2) = c_B^2(s_B^1, s_B^2) = s_B^1 + s_B^2$. Note that signal component s_B^1 of supplier 1 does not affect his cost in either of the two possible outcomes. Incentive compatibility therefore requires that the probability of supplier 1 being chosen to produce specification A is independent of s_B^1. This, however, runs counter to the goal of achieving efficiency, and hence, there exists no efficient, incentive compatible DRM.

[7] Jehiel and Moldovanu (2001) derive a generalization of Myerson's (1981) "revenue equivalence" to the case of multi-dimensional dimensional types.

4.6 Appendix 1: Incentive Compatible Procurement Mechanisms

In this section we take a general mechanism design approach to the design of procurement procedures. It forms the foundation for any future work on expenditure minimizing procurement mechanisms. We show that, in contrast to the design of optimal auctions with one-dimensional private information, one cannot characterize an optimal (i.e. expenditure minimizing) DRM by (initially) neglecting the incentive compatibility constraints. Recall that in the case of one-dimensional signals, incentive compatibility is equivalent to the requirement that the expected conditional allocation probabilities are nondecreasing.[8] We find that this approach does not work with multi-dimensional signals, because a solution to the buyer's unconstrained expenditure minimization problem will typically not satisfy the constraints imposed by incentive compatibility. We therefore need to find a way to explicitly account for the incentive compatibility constraints in the buyer's optimization problem. This is left for future research.

When looking for expenditure-minimizing mechanisms, the Revelation Principle tells us that we can restrict attention to direct revelation mechanisms (DRM) under which the suppliers report their signals in exchange for a (report-contingent) payment. Given reports s_A^1, s_B^1, s_A^2 and s_B^2 a DRM specifies:

- probability $\pi_{Ai}(s_A^1, s_B^1, s_A^2, s_B^2)$ that supplier i ($i = 1, 2$) is chosen to produce specification A
- probability $\pi_{Bi}(s_A^1, s_B^1, s_A^2, s_B^2)$ that supplier i is chosen to produce specification B
- payment $\tau_i(s_A^i, s_B^i)$ to supplier i for revealing his signals

There is no loss of generality in restricting ourselves to payments τ_i that depend only on supplier i's signals, and not on those of supplier j. The reason is that even if

[8] See Krishna (2002), p.68.

the payment function τ_i depends on supplier j's signals, supplier i's decision whether or not to reveal his signals truthfully is made only on the basis of the expected value of τ_i.

In the following, we only consider supplier 1. The expressions for supplier 2 are completely analogous. We are looking for restrictions on the DRM which ensure that supplier 1 truthfully reveals his signals if supplier 2 does so. Supplier 1's expected payoff when he reports his signal vector as $\widehat{\mathbf{s}}^1 = (\widehat{s}_A^1, \widehat{s}_B^1)$ although the true signal vector is $\mathbf{s}^1$ is given by[9]

$$U_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1) := \tau_1(\widehat{\mathbf{s}}^1) - \int_{[0,1]^2} \pi_{A1}(\widehat{\mathbf{s}}^1, \mathbf{s}^2)(s_A^1 + s_A^2) f(\mathbf{s}^2) d\mathbf{s}^2 \\ - \int_{[0,1]^2} \pi_{B1}(\widehat{\mathbf{s}}^1, \mathbf{s}^2)(s_B^1 + s_B^2) f(\mathbf{s}^2) d\mathbf{s}^2.$$

Supplier 1's decision problem is the following: Given the true signal vector $\mathbf{s}^1$, choose the report $\widehat{\mathbf{s}}^1 \in [0,1]^2$ that solves:

$$\max_{\widehat{\mathbf{s}}^1} U_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1).$$

The following notation simplifies the analysis: Define supplier 1's expected conditional probabilities of being chosen to produce specification $k \in \{A, B\}$ as follows:

$$q_{k1}(\widehat{\mathbf{s}}^1) := \int_{[0,1]^2} \pi_{k1}(\widehat{\mathbf{s}}^1, \mathbf{s}^2) f(\mathbf{s}^2) d\mathbf{s}^2.$$

Supplier 1's expected production cost when participating in the DRM can therefore be written as follows:

$$\begin{aligned} C_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1) &= s_A^1 q_{A1}(\widehat{\mathbf{s}}^1) + s_B^1 q_{B1}(\widehat{\mathbf{s}}^1) \\ &+ \int_{[0,1]^2} \pi_{A1}(\widehat{\mathbf{s}}^1, \mathbf{s}^2) s_A^2 f(\mathbf{s}^2) d\mathbf{s}^2 \\ &+ \int_{[0,1]^2} \pi_{B1}(\widehat{\mathbf{s}}^1, \mathbf{s}^2) s_B^2 f(\mathbf{s}^2) d\mathbf{s}^2. \end{aligned}$$

Thus, supplier 1's expected payoff when his signal vector is $\mathbf{s}^1$ and he reports it as $\widehat{\mathbf{s}}^1$ can be written as:

$$U_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1) = \tau_1(\widehat{\mathbf{s}}^1) - C_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1).$$

[9]To easy notation, define the joint density $f(s_A^i, s_B^i) = g(s_A^i) g(s_B^i)$ for all $i = 1, 2$.

Define supplier 1's expected payoff/rent from truthful revelation (i.e. $\widehat{\mathbf{s}}^1 = \mathbf{s}^1$) as follows:

$$\mu_1(\mathbf{s}^1) := U_1(\mathbf{s}^1, \mathbf{s}^1) = \tau_1(\mathbf{s}^1) - C_1(\mathbf{s}^1, \mathbf{s}^1).$$

A DRM is said to be incentive compatible if truthful revelation maximizes supplier 1's expected rents. That is,

$$\mu_1(\mathbf{s}^1) = \max_{\widehat{\mathbf{s}}^1} \left(\tau_1(\widehat{\mathbf{s}}^1) - C_1(\widehat{\mathbf{s}}^1, \mathbf{s}^1)\right).$$

As participation in the DRM is voluntary for supplier 1, the participation constraint is that $\mu_1(\mathbf{s}^1) \geq 0$. Note that supplier 1's expected rents from truthful revelation, μ_1, is a convex function, and is therefore differentiable almost everywhere in $[0,1]^2$. It follows via the Envelope Theorem that

$$\frac{\partial \mu_1(\mathbf{s}^1)}{\partial s_A^1} = -q_{A1}(\mathbf{s}^1)$$

and

$$\frac{\partial \mu_1(\mathbf{s}^1)}{\partial s_B^1} = -q_{B1}(\mathbf{s}^1).$$

As q_{A1} and q_{B1} are probabilities, it follows immediately that $\mu_1(\mathbf{s}^1)$ decreasing on $[0,1]^2$. Note that monotonicity of μ_1 reduces the participation constraint to $\mu_1(1,1) \geq 0$. Thus, as we assume that the buyer has to ensure provision of the good and hence participation, it is optimal to set $\mu_1(1,1) = 0$. Note furthermore that the convexity of μ_1 implies that

$$\frac{q_{k1}(\mathbf{s}^1)}{\partial s_k^1} \leq 0 \text{ for all } k \in \{A, B\}.$$

The expected rents of supplier 1 from participating in an incentive compatible DRM depend only on the (conditional expected) allocation probabilities. The above envelope condition implies that

$$\nabla \mu_1(\mathbf{s}^1) = \left\langle -q_{A1}(\mathbf{s}^1), -q_{B1}(\mathbf{s}^1)\right\rangle. \tag{4.9}$$

Any real-valued and differentiable function μ_1 on $[0,1]^2$ such that $\nabla \mu_1(\mathbf{s}^1) = \langle -q_{A1}(\mathbf{s}^1), -q_{B1}(\mathbf{s}^1)\rangle$ is called *potential function.* Thus, incentive compatibility implies that μ_1 is a potential function for the vector field $\langle -q_{A1}(\mathbf{s}^1), -q_{B1}(\mathbf{s}^1)\rangle$. Note that a necessary condition for the existence of a potential function μ_1 is that

$$\frac{\partial q_{A1}(s_A^1, s_B^1)}{\partial s_B^1} = \frac{\partial q_{B1}(s_A^1, s_B^1)}{\partial s_A^1}. \tag{4.10}$$

The obvious question is how to find potential function μ_1 from (4.9)? Some multivariate version of the "Fundamental Theorem of Calculus" regarding the relationship between integration and differentiation is needed.[10] Note that if the necessary condition in (4.10) is satisfied, then we can express $\mu_1(s_A^1, s_B^1)$ as the integral of $\nabla\mu_1(s_A^1, s_B^1)$ along *any curve* in $[0,1]^2$ that connects the points (s_A^1, s_B^1) and $(1,1)$. Such a curve $\boldsymbol{\sigma}^1(r) = \langle \sigma_A^1(r), \sigma_B^1(r)\rangle$ is parameterized by $r \in [0,1]$ so that

$$\begin{aligned} \langle \sigma_A^1(0), \sigma_B^1(0)\rangle &= (1,1) \\ \langle \sigma_A^1(1), \sigma_B^1(1)\rangle &= (s_A^1, s_B^1). \end{aligned}$$

The curve integral of μ_1 between the points (s_A^1, s_B^1) and $(1,1)$ is defined as follows:

$$\begin{aligned} \mu_1(s_A^1, s_B^1) &= \int_0^1 \frac{d}{dr}\mu_1(\sigma_A^1(r), \sigma_B^1(r))dr \\ &= \mu_1(\sigma_A^1(1), \sigma_B^1(1)) - \mu_1(\sigma_A^1(0), \sigma_B^1(0)) \\ &= \mu_1(s_A^1, s_B^1) - \mu_1(1,1) = \mu_1(s_A^1, s_B^1). \end{aligned}$$

Note that $\frac{d}{dr}\mu_1(\sigma_A^1(r), \sigma_B^1(r))$ is given by

$$\frac{\partial \mu_1(\boldsymbol{\sigma}^1(r))}{\partial s_A^1}\frac{d\sigma_A^1(r)}{dr} + \frac{\partial \mu_1(\boldsymbol{\sigma}^1(r))}{\partial s_B^1}\frac{d\sigma_B^1(r)}{dr}.$$

Thus, for an arbitrary path $\boldsymbol{\sigma}^1(r)$ we can express $\mu_1$as follows:

$$-\int_0^1 \left(q_{A1}(\boldsymbol{\sigma}^1(r))\frac{d\sigma_A^1(r)}{dr} + q_{B1}(\boldsymbol{\sigma}^1(r))\frac{d\sigma_B^1(r)}{dr} \right) dr.$$

The next question is which path one should choose? One possibility is to integrate $\nabla\mu_1(s_A^1, s_B^1)$ along the *line* connecting the points (s_A^1, s_B^1) and $(1,1)$. That is,

$$\boldsymbol{\sigma}^1(r) = \langle 1 - r(1 - s_A^1), 1 - r(1 - s_B^1)\rangle.$$

[10]See Lang (1998).

In this case, $\mu_1(s_A^1, s_B^1)$ is given by

$$\int_0^1 \Big(q_{A1}(1-r(1-s_A^1), 1-r(1-s_B^1))(1-s_A^1) + q_{B1}(1-r(1-s_A^1), 1-r(1-s_B^1))(1-s_A^1)\Big)\, dr.$$

We have now fully expressed μ_1 in terms of the (conditional expected) allocation probabilities q_{A1} and q_{B1}. Thus, supplier 1's payment τ_1 is fully determined:

$$\tau_1(\mathbf{s}^1) = \mu_1(\mathbf{s}^1) + C_1(\mathbf{s}^1, \mathbf{s}^1).$$

The problem of designing an expenditure minimizing DRM is now to solve

$$\min_{\tau_1, \tau_2} \int_{[0,1]^2} \tau_1(\mathbf{s}^1) f(\mathbf{s}^1) d\mathbf{s}^1 + \int_{[0,1]^2} \tau_2(\mathbf{s}^2) f(\mathbf{s}^2) d\mathbf{s}^2.$$

Note that the expected payment to supplier 1 under an incentive compatible DRM is given by

$$E_{\mathbf{s}^1}[\tau_1(\mathbf{s}^1)] = E_{\mathbf{s}^1}[\mu_1(\mathbf{s}^1) + C_1(\mathbf{s}^1, \mathbf{s}^1)].$$

In order to express the expected rents $E_{\mathbf{s}^1}[\mu_1(\mathbf{s}^1)]$ in a more compact way, we can follow a technique used by Armstrong (1996), which is a form of "multivariate integration by parts". It yields the following expression:

$$E_{\mathbf{s}^1}[\mu_1(\mathbf{s}^1)] = \int_{[0,1]^2} \left(q_{A1}(\mathbf{s}^1)(1-s_A^1) + q_{B1}(\mathbf{s}^1)(1-s_B^1)\right) h(\mathbf{s}^1) d\mathbf{s}^1,$$

where $h(s^1)$ is given by

$$\int_1^{\min\{\frac{1}{1-s_A^1}, \frac{1}{1-s_B^1}\}} r f(1-r(1-s_A^1), 1-r(1-s_B^1)) dr.$$

We can therefore express $E_{\mathbf{s}^1}[\mu_1(\mathbf{s}^1) + C_1(\mathbf{s}^1, \mathbf{s}^1)]$ as

$$E_{\mathbf{s}^1}\left[\sum_{k\in\{A,B\}} \left(q_{k1}(s^1)\left(s_k^1 + (1-s_k^1)\frac{h(s^1)}{f(s^1)}\right)\right) + \sum_{k\in\{A,B\}} \left(\int_{[0,1]^2} \pi_{k1}(s^1, s^2) s_k^2 f(s^2) ds^2\right)\right].$$

The buyer's problem is therefore to choose $\pi_{A1}, \pi_{B1}, \pi_{A2}$ and π_{B2} so as to minimize the sum of the suppliers' expected payments. We can express supplier 1's expected payment in terms of π_{A1}, π_{B1}:

$$E_{\mathbf{s}^1,\mathbf{s}^2}\left[\sum_{k\in\{A,B\}}\left(\pi_{k1}(s^1,s^2)\left(s_k^1+s_k^2+(1-s_k^1)\frac{h(s^1)}{f(s^1)}\right)\right)\right]$$

To obtain the buyer's objective function, simply add up the expected payments to suppliers 1 and 2. This yields:

$$E_{\mathbf{s}^1,\mathbf{s}^2}\left[\sum_{k\in\{A,B\}}\left(\pi_{k1}(s^1,s^2)\left(s_k^1+s_k^2+(1-s_k^1)\frac{h(s^1)}{f(s^1)}\right)\right.\right.$$
$$\left.\left.+\pi_{k2}(s^1,s^2)\left(s_k^1+s_k^2+(1-s_k^2)\frac{h(s^2)}{f(s^2)}\right)\right)\right].$$

In the uniform case, the buyer's problem is to minimize the expected value of following expression:

$$\begin{aligned}
&\pi_{A1}(s^1,s^2)\left(s_A^1+s_A^2+(1-s_A^1)\frac{1-(1-\min\{s_A^1,s_B^1\})^2}{2(1-\min\{s_A^1,s_B^1\})^2}\right)\\
&+\pi_{B1}(s^1,s^2)\left(s_B^1+s_B^2+(1-s_B^1)\frac{1-(1-\min\{s_A^1,s_B^1\})^2}{2(1-\min\{s_A^1,s_B^1\})^2}\right)\\
&+\pi_{A2}(s^1,s^2)\left(s_A^1+s_A^2+(1-s_A^2)\frac{1-(1-\min\{s_A^2,s_B^2\})^2}{2(1-\min\{s_A^2,s_B^2\})^2}\right)\\
&+\pi_{B2}(s^1,s^2)\left(s_B^1+s_B^2+(1-s_B^2)\frac{1-(1-\min\{s_A^2,s_B^2\})^2}{2(1-\min\{s_A^2,s_B^2\})^2}\right),
\end{aligned}$$

subject to the constraint that the π-functions must be chosen such that the necessary condition for the existence of potential functions μ_1, μ_2 is satisfied:

$$\frac{\partial q_{Ai}(s_A^i,s_B^i)}{\partial s_B^i}=\frac{\partial q_{Bi}(s_A^i,s_B^i)}{\partial s_A^i}.$$

4.7 Appendix 2: Proofs

Proof of Lemma 4.1: Suppose the suppliers' strategies $(k^{1*},p_A^{1*},p_B^{1*})$ and $(k^{2*},p_A^{2*},p_B^{2*})$ constitute a BNE of the alternative version of the MPM. This implies that for every signal-vector $\mathbf{s}^1$, the strategy $(k^{1*},p_A^{1*},p_B^{1*})$ is a best response to supplier 2's strategy

$(k^{2*}, p_A^{2*}, p_B^{2*})$ (and vice versa). In particular, given supplier 1's equilibrium specification choice k^{1*}, the pricing functions p_A^{1*} and p_B^{1*} solve

$$\max_{p_A^1, p_B^1} E_{\mathbf{s}^2}[(p^1_{k^{1*}(\mathbf{s}^1)}(\mathbf{s}^1) - c_{k^{1*}(\mathbf{s}^1)})] \Pr\{p^{2*}_{k^{2*}(\mathbf{s}^2)}(\mathbf{s}^2) > p^1_{k^{1*}(\mathbf{s}^1)}(\mathbf{s}^1)\}, \tag{4.11}$$

for every $\mathbf{s}^1 \in S$. Now consider the original version of the MPM. Supplier 1's strategy consists of the pricing functions q_A^1 and q_B^1, and his expected payoff is

$$E_{\mathbf{s}^2}[(q^1_{l^1(\mathbf{s}^1)}(\mathbf{s}^1) - c_{l^1(\mathbf{s}^1)})] \Pr\{\min\{q_A^2(\mathbf{s}^2), q_B^2(\mathbf{s}^2)\} > q^1_{l^1(\mathbf{s}^1)}(\mathbf{s}^1)\}, \tag{4.12}$$

where $l^1(\mathbf{s}^1) = \arg\min_{l \in \{A,B\}} q_l^1(\mathbf{s}^1)$. Based on his equilibrium strategy in the alternative MPM, we can construct for each supplier i ($i = 1, 2$) the following strategy for the original MPM:

$$\widehat{q}_A^i(\mathbf{s}^i) = \begin{cases} p_A^{i*}(\mathbf{s}^i) & \text{if } k^{i*}(\mathbf{s}^i) = A \\ \max_{\mathbf{t}^i \in S} p_B^{i*}(\mathbf{t}^i) & \text{if } k^{i*}(\mathbf{s}^i) = B \end{cases} \tag{4.13}$$

and

$$\widehat{q}_B^i(\mathbf{s}^i) = \begin{cases} \max_{\mathbf{t}^i \in S} p_A^{i*}(\mathbf{t}^i) & \text{if } k^{i*}(\mathbf{s}^i) = A \\ p_B^{i*}(\mathbf{s}^i) & \text{if } k^{i*}(\mathbf{s}^i) = B \end{cases}. \tag{4.14}$$

The interpretation is as follows: The pricing function $\widehat{q}_A^i$ equals the equilibrium pricing function p_A^{i*} for specification A if supplier i's signal $\mathbf{s}^i$ is such that he would choose specification A in the alternative MPM. Otherwise, it takes on a fixed value equal to $\max_{\mathbf{t}^i \in S} p_B^{i*}(\mathbf{t}^i)$, which is the highest possible price supplier i would quote for specification B under the alternative MPM. Suppose now that supplier 2 uses strategy $(\widehat{q}_A^2, \widehat{q}_B^2)$ in the original MPM. It follows that

$$\min\{\widehat{q}_A^2(\mathbf{s}^j), \widehat{q}_B^2(\mathbf{s}^j)\} = \begin{cases} p_A^{2*}(\mathbf{s}^2) & \text{if } k^{2*}(\mathbf{s}^2) = A \\ p_B^{2*}(\mathbf{s}^2) & \text{if } k^{2*}(\mathbf{s}^2) = B \end{cases},$$

or, more compactly,

$$\min\{\widehat{q}_A^2(\mathbf{s}^j), \widehat{q}_B^2(\mathbf{s}^j)\} = p^{2*}_{k^{2*}(\mathbf{s}^2)}(\mathbf{s}^2).$$

We can therefore write supplier 1's expected payoff in (4.12) as follows:

$$E_{\mathbf{s}^2}[(q^1_{l^1(\mathbf{s}^1)}(\mathbf{s}^1) - c_{l^1(\mathbf{s}^1)})] \Pr\{p^{2*}_{k^{2*}(\mathbf{s}^2)}(\mathbf{s}^2) > q^1_{l^1(\mathbf{s}^1)}(\mathbf{s}^1)\}. \tag{4.15}$$

It is easy to see that the strategy $(\widehat{q}_A^1, \widehat{q}_B^1)$ given by (4.13) and (4.14) constitutes a best response for supplier 1 as the pricing functions p_A^{1*} and p_B^{1*} solve

$$\max_{p_A^1, p_B^1} E_{\mathbf{s}^2}[(p^1_{k^{1*}(\mathbf{s}^1)}(\mathbf{s}^1) - c_{k^{1*}(\mathbf{s}^1)})] \Pr\{p^{2*}_{k^{2*}(\mathbf{s}^2)}(\mathbf{s}^2) > p^1_{k^{1*}(\mathbf{s}^1)}(\mathbf{s}^1)\}.$$

This shows that every BNE of the alternative MPM constitutes a BNE of the original MPM.

Q.E.D.

Proof of Proposition 4.2: If supplier 1 with signal (s_A^1, s_B^1) (where $0 < s_B^1 < s_A^1 < 1$) reports the signal values $(\widehat{s}_A^1, \widehat{s}_B^1)$, his expected profit is

$$U_1(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1) := \tau_1^{eff} - C_1(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1),$$

where

$$\begin{aligned} C_1(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1) := {} & s_A^1 \int_S\int (\pi_{A1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s_B^2) g(s_A^2) g(s_B^2) ds_A^2 ds_B^2 \\ & + s_B^1 \int_S\int \pi_{B1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s_B^2) g(s_A^2) g(s_B^2) ds_A^2 ds_B^2 \\ & + \int_S\int (\pi_{A1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s_B^2) s_A^2 g(s_A^2) g(s_B^2) ds_A^2 ds_B^2 \\ & + \int_S\int \pi_{B1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s_B^2) s_B^2 g(s_A^2) g(s_B^2) ds_A^2 ds_B^2. \end{aligned}$$

In the following, we assume initially that the probability with which supplier 1 is chosen to produce the efficient specification is a function of the suppliers' reports. That is,

$$\pi_{A1}^{eff}(s_A^1, s_B^1, s_A^2, s_B^2) = \begin{cases} \lambda_1(s_A^1, s_B^1, s_A^2, s_B^2) & \text{if } s_A^1 + s_A^2 < s_B^1 + s_B^2 \\ 0 & \text{if } s_A^1 + s_A^2 > s_B^1 + s_B^2 \end{cases}$$

and

$$\pi_{B1}^{eff}(s_A^1, s_B^1, s_A^2, s_B^2) = \begin{cases} 0 & \text{if } s_A^1 + s_A^2 < s_B^1 + s_B^2 \\ \lambda_1(s_A^1, s_B^1, s_A^2, s_B^2) & \text{if } s_A^1 + s_A^2 > s_B^1 + s_B^2 \end{cases}$$

Now consider the expected probability with which supplier 1 is chosen to produce specification A if his signals are s_A^1 and s_B^1 but he reports them as $\widehat{s}_A^1$ and $\widehat{s}_B^1$:

$$\begin{aligned} & \int_S\int (\pi_{A1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s_B^2) g(s_A^2) g(s_B^2) ds_A^2 ds_B^2 \\ & \qquad = \int_0^{1-\widehat{s}_A^1+\widehat{s}_B^1} \left(\int_{s_A^2+\widehat{s}_A^1-\widehat{s}_B^1}^{1} \lambda_1(s_A^1, s_B^1, s_A^2, s_B^2) g(s_B^2) ds_B^2 \right) g(s_A^2) ds_A^2 \end{aligned}$$

and

$$\int_S\int \pi^{eff}_{B1}(\widehat{s}^1_A,\widehat{s}^1_B,s^2_A,s^2_B)g(s^2_A)g(s^2_B)ds^2_Ads^2_B$$
$$=\int_0^{1-\widehat{s}^1_A+\widehat{s}^1_B}\left(\int_0^{s^2_A+\widehat{s}^1_A-\widehat{s}^1_B}\lambda_1(s^1_A,s^1_B,s^2_A,s^2_B)g(s^2_B)ds^2_B\right)g(s^2_A)ds^2_A$$
$$+\int_{1-\widehat{s}^1_A+\widehat{s}^1_B}^1\left(\int_0^1\lambda_1(s^1_A,s^1_B,s^2_A,s^2_B)g(s^2_B)ds^2_B\right)g(s^2_A)ds^2_A.$$

Jehiel and Moldovanu (2001) show that a necessary and sufficient condition for truthful revelation of supplier 1's signals is

$$\frac{\partial}{\partial\widehat{s}^1_B}\int_S\int(\pi^{eff}_{A1}(\widehat{s}^1_A,\widehat{s}^1_B,s^2_A,s^2_B)g(s^2_A)g(s^2_B)ds^2_Ads^2_B$$
$$=\frac{\partial}{\partial\widehat{s}^1_A}\int_S\int\pi^{eff}_{B1}(\widehat{s}^1_A,\widehat{s}^1_B,s^2_A,s^2_B)g(s^2_A)g(s^2_B)ds^2_Ads^2_B.$$

With the above specification of π^{eff}_{A1} and π^{eff}_{B1} we obtain:

$$\frac{\partial}{\partial\widehat{s}^1_B}\int_S\int(\pi^{eff}_{A1}(\widehat{s}^1_A,\widehat{s}^1_B,s^2_A,s^2_B)g(s^2_A)g(s^2_B)ds^2_Ads^2_B$$
$$=\int_0^{1-\widehat{s}^1_A+\widehat{s}^1_B}\left(\int_{s^2_A+\widehat{s}^1_A-\widehat{s}^1_B}^1\frac{\partial\lambda_1}{\partial\widehat{s}^1_B}g(s^2_B)ds^2_B\right)g(s^2_A)ds^2_A$$
$$+\int_0^{1-\widehat{s}^1_A+\widehat{s}^1_B}\lambda_1(\widehat{s}^1_A,\widehat{s}^1_B,s^2_A,s^2_A+\widehat{s}^1_A-\widehat{s}^1_B)g(s^2_A+\widehat{s}^1_A-\widehat{s}^1_B)g(s^2_A)ds^2_A$$

and

$$\frac{\partial}{\partial\widehat{s}^1_A}\int_S\int\pi^{eff}_{B1}(\widehat{s}^1_A,\widehat{s}^1_B,s^2_A,s^2_B)g(s^2_A)g(s^2_B)ds^2_Ads^2_B$$
$$=\int_0^{1-\widehat{s}^1_A+\widehat{s}^1_B}\left(\int_0^{s^2_A+\widehat{s}^1_A-\widehat{s}^1_B}\frac{\partial\lambda_1}{\partial\widehat{s}^1_A}g(s^2_B)ds^2_B\right)g(s^2_A)ds^2_A$$
$$+\int_{1-\widehat{s}^1_A+\widehat{s}^1_B}^1\left(\int_0^1\frac{\partial\lambda_1}{\partial\widehat{s}^1_A}g(s^2_B)ds^2_B\right)g(s^2_A)ds^2_A$$
$$+\int_0^{1-\widehat{s}^1_A+\widehat{s}^1_B}\lambda_1(\widehat{s}^1_A,\widehat{s}^1_B,s^2_A,s^2_A+\widehat{s}^1_A-\widehat{s}^1_B)g(s^2_A+\widehat{s}^1_A-\widehat{s}^1_B)g(s^2_A)ds^2_A.$$

It is easy to see that the necessary and sufficient condition in Jehiel and Moldovanu (2001) is satisfied if λ_1 is a constant. In the following we therefore restrict attention to this case. For ease of notation let $\lambda_1 = \frac{1}{2}$. This yields

$$\int_S\int(\pi_{A1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s_B^2)g(s_A^2)g(s_B^2)ds_A^2 ds_B^2 = \frac{1}{2}\int_0^{1-\widehat{s}_A^1+\widehat{s}_B^1}\left(\int_{s_A^2+\widehat{s}_A^1-\widehat{s}_B^1}^1 g(s_B^2)ds_B^2\right)g(s_A^2)ds_A^2$$

and

$$\int_S\int \pi_{B1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s_B^2)g(s_A^2)g(s_B^2)ds_A^2 ds_B^2 = \frac{1}{2}\int_0^{1-\widehat{s}_A^1+\widehat{s}_B^1}\left(\int_0^{s_A^2+\widehat{s}_A^1-\widehat{s}_B^1} g(s_B^2)ds_B^2\right)g(s_A^2)ds_A^2$$
$$+\frac{1}{2}\int_{1-\widehat{s}_A^1+\widehat{s}_B^1}^1\left(\int_0^1 g(s_B^2)ds_B^2\right)g(s_A^2)ds_A^2.$$

Also,

$$\int_0^1\int_0^1 \pi_{A1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s)s_A^2 g(s_A^2)g(s_B^2)ds_A^2 ds_B^2 = \int_0^{1-(\widehat{s}_A^1-\widehat{s}_B^1)}\left(\int_{s_A^2+(\widehat{s}_A^1-\widehat{s}_B^1)}^1 \frac{1}{2}s_A^2 g(s_B^2)ds_B^2\right)g(s_A^2)($$

and

$$\int_0^1\int_0^1 \pi_{B1}^{eff}(\widehat{s}_A^1, \widehat{s}_B^1, s_A^2, s_B^2)s_B^2 g(s_A^2)g(s_B^2)ds_A^2 ds_B^2 = \int_0^{1-(\widehat{s}_A^1-\widehat{s}_B^1)}\left(\int_0^{s_A^2+(\widehat{s}_A^1-\widehat{s}_B^1)} \frac{1}{2}s_B^2 g(s_B^2)ds_B^2\right)g(s_A^2)d$$
$$+\int_{1-(\widehat{s}_A^1-\widehat{s}_B^1)}^1\left(\int_0^1 \frac{1}{2}s_B^2 g(s_B^2)ds_B^2\right)g(s_A^2)ds_A^2.$$

It is then easy to compute the gradient vector $DC(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)$:

$$\frac{\partial C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{\partial \widehat{s}_A^1} = -\frac{1}{2}((s_A^1 - \widehat{s}_A^1) - (s_B^1 - \widehat{s}_B^1))\int_0^{1-\widehat{s}_A^1+\widehat{s}_B^1} g(s_A^2 + \widehat{s}_A^1 - \widehat{s}_B^1)g(s_A^2)ds_A^2$$

and

$$\frac{\partial C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{\partial \widehat{s}_B^1} = -\frac{\partial C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{\partial \widehat{s}_A^1}.$$

Thus, $DC(s_A^1, s_B^1, s_A^1, s_B^1) = \mathbf{0}$. As τ^{eff} is constant we therefore have $DU(s_A^1, s_B^1, s_A^1, s_B^1) = \mathbf{0}$, so that the necessary condition for a maximum of $U(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)$ at the point $(\widehat{s}_A^1, \widehat{s}_B^1) = (s_A^1, s_B^1)$ is satisfied.

Now compute the Hessian matrix $D^2C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)$. Its elements are as follows:

$$\frac{\partial^2 C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{\partial \widehat{s}_A^{12}} = \frac{1}{2}\int_0^{1-\widehat{s}_A^1+\widehat{s}_B^1} g(s_A^2 + \widehat{s}_A^1 - \widehat{s}_B^1) g(s_A^2) ds_A^2$$
$$-\frac{1}{2}((s_A^1 - \widehat{s}_A^1) - (s_B^1 - \widehat{s}_B^1)) \left(\int_0^{1-\widehat{s}_A^1+\widehat{s}_B^1} g'(s_A^2 + \widehat{s}_A^1 - \widehat{s}_B^1) g(s_A^2) ds_A^2 - g(1) g(1 - \widehat{s}_A^1 + \widehat{s}_B^1) \right),$$

$$\frac{\partial^2 C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{\partial \widehat{s}_A^1 \partial \widehat{s}_B^1} = -\frac{\partial^2 C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{(\partial \widehat{s}_A^1)^2}$$

and

$$\frac{\partial^2 C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{(\partial \widehat{s}_B^1)^2} = -\frac{\partial^2 C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{\partial \widehat{s}_A^1 \partial \widehat{s}_B^1} = \frac{\partial^2 C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)}{(\partial \widehat{s}_A^1)^2}.$$

In the uniform case, for example, the Hessian matrix $D^2C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)$ reduces to:

$$\begin{pmatrix} -\widehat{s}_A^1 - \frac{1}{2}s_B^1 + \widehat{s}_B^1 + \frac{1}{2}s_A^1 + \frac{1}{2} & \widehat{s}_A^1 + \frac{1}{2}s_B^1 - \widehat{s}_B^1 - \frac{1}{2}s_A^1 - \frac{1}{2} \\ \widehat{s}_A^1 + \frac{1}{2}s_B^1 - \widehat{s}_B^1 - \frac{1}{2}s_A^1 - \frac{1}{2} & -\widehat{s}_A^1 - \frac{1}{2}s_B^1 + \widehat{s}_B^1 + \frac{1}{2}s_A^1 + \frac{1}{2} \end{pmatrix}.$$

Evaluating the general Hessian at $(\widehat{s}_A^1, \widehat{s}_B^1) = (s_A^1, s_B^1)$ we obtain $D^2C(s_A^1, s_B^1, s_A^1, s_B^1)$:

$$\begin{pmatrix} \frac{1}{2}\int_0^{1-s_A^1+s_B^1} g(s_A^2 + s_A^1 - s_B^1) g(s_A^2) ds_A^2 & -\frac{1}{2}\int_0^{1-s_A^1+s_B^1} g(s_A^2 + s_A^1 - s_B^1) g(s_A^2) ds_A^2 \\ -\frac{1}{2}\int_0^{1-s_A^1+s_B^1} g(s_A^2 + s_A^1 - s_B^1) g(s_A^2) ds_A^2 & \frac{1}{2}\int_0^{1-s_A^1+s_B^1} g(s_A^2 + s_A^1 - s_B^1) g(s_A^2) ds_A^2 \end{pmatrix},$$

which is a positive *semi*definite matrix as $\det(D^2C(s_A^1, s_B^1, s_A^1, s_B^1)) = 0$ (ambiguous case). However, note that the function $C(\widehat{s}_A^1, s_B^1, s_A^1, s_B^1)$ has a minimum at $\widehat{s}_A^1 = s_A^1$ (at this point, the second derivative is $\frac{1}{2}\int_0^{1-s_A^1+s_B^1} g(s_A^2 + s_A^1 - s_B^1) g(s_A^2) ds_A^2$, which is strictly positive for all $s_A^1, s_B^1 \in (0,1)$ and $s_A^1 > s_B^1$). Likewise, the function $C(s_A^1, \widehat{s}_B^1, s_A^1, s_B^1)$ has a minimum at $\widehat{s}_B^1 = s_B^1$. Thus, we can conclude that $C(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)$ has a minimum at $(\widehat{s}_A^1, \widehat{s}_B^1) = (s_A^1, s_B^1)$.

As the transfer function τ^{eff} is constant, the expected profit $U(\widehat{s}_A^1, \widehat{s}_B^1, s_A^1, s_B^1)$ is maximal at $(\widehat{s}_A^1, \widehat{s}_B^1) = (s_A^1, s_B^1)$. An analogous argument can be made for all signal values $s_A^1, s_B^1 \in (0,1)$ where $s_A^1 \leq s_B^1$, and hence the efficient DRM is incentive compatible.

Q.E.D.

Chapter 5

Concluding Remarks

In this thesis we have adopted a mechanism design approach to study three specific situations of economic relevance. We have thereby added a new angle to important themes that , in one form or another, have been at the heart of the mechanism design literature.

Chapter 2 on incentive contracts for information providers emphasizes that sequential information acquisition procedures that rely on a direct comparison of the payoffs of two alternatives cannot be costlessly delegated to a self-interested agent, even in our setup that is more favorable to the principal as one would ever expect to find in reality.

Chapter 3 studied whether efficient procedures exist that give agents with conflicting interests the incentive to truthfully reveal their "strength" of preference when money transfers are ruled out. Our main result was that no such mechanisms exist if the distribution of agents' types has continuum support. If three types have positive prior probability, there exist examples in which incentive compatible and ex post efficient decision rules exist, and others where they do not. We have numerically characterized a second best decision rule for the latter case numerically. It is highly sensitive to the model parameters, and hence an important question for future research is the characterization of "robust" rules for compromising.

Finally, chapter 4 has established in a simple procurement model for a good that exists in two specifications, that when suppliers have multi-dimensional private information an optimal procurement mechanism will not generally be efficient. A full characterization of optimal mechanisms in the presence of multi-dimensional private information is a very difficult problem. General results to date exist only for multi-dimensional monopoly screening.[1]

[1]See Rochet and Choné (1998).

Bibliography

[1] ARMSTRONG, MARK (1996), "Multiproduct Nonlinear Pricing", *Econometrica*, Vol. 64, No. 1, pp. 51-75.

[2] BERGEMANN, DIRK AND STEPHEN MORRIS (2003), "Robust Mechanism Design", Cowles Foundation, Yale University, mimeo.

[3] BÖRGERS, TILMAN (1991), "Undominated Strategies and Coordination in Normalform Games", *Social Choice and Welfare*, Vol. 8, pp. 65-78.

[4] BRAMS, STEVEN J. AND PETER C. FISHBURN (1983), *Approval Voting*, Boston: Birkhauser.

[5] CRÉMER, JACQUES AND FAHAD KHALIL (1992), "Gathering Information before Signing a Contract", *American Economic Review*, Vol. 82 (3), pp. 566-578.

[6] CRÉMER, JACQUES, FAHAD KHALIL AND JEAN-CHARLES ROCHET (1998a), "Strategic Information Gathering before a Contract is Offered", *Journal of Economic Theory*, Vol. 81, pp. 163-200.

[7] CRÉMER, JACQUES, FAHAD KHALIL AND JEAN-CHARLES ROCHET (1998b), "Contracts and productive information gathering", *Games and Economic Behavior*, Vol. 25, pp. 174-193

[8] GUESNERIE, ROGER, PIERRE PICARD AND PATRICK REY (1998), "Adverse selection and moral hazard with risk-neutral agents", *European Economic Review*, Vol. 33, pp. 807-823

[9] HART, OLIVER AND BENGT HOLMSTRÖM, "The theory of contracts", in T.F. Bewley, ed., *Advances in Economic Theory: Fifth World Congress*, Cambridge: Cambridge University Press, chapter 3, pp. 71-155

[10] HOLMSTRÖM, BENGT, AND ROGER MYERSON (1983), "Efficient and Durable Decision Rules with Incomplete Information", *Econometrica* Vol. 51, pp. 1799-1820.

[11] JEHIEL, PHILIPPE AND BENNY MOLDOVANU (2001), "Efficient Design with Interdependent Valuations", *Econometrica*, Vol. 69, No.5, pp. 1237-1259.

[12] KRISHNA, VIJAY (2002), *Auction Theory*, San Diego: Academic Press.

[13] KRISHNA, VIJAY AND MOTTY PERRY (2000), "Efficient Mechanism Design", Pennsylvania State University mimeo.

[14] LANG, SERGE (1988): *Calculus of Several Variables*, Heidelberg: Spinger Verlag.

[15] MALCOMSON, JAMES M. (2002), "Principal and Expert Agent", University of Oxford mimeo.

[16] MAS-COLELL, ANDREU, MICHAEL D. WHINSTON AND JERRY R. GREEN (1995), *Microeconomic Theory*, Oxford: Oxford University Press.

[17] MASKIN, ERIC S. (1992), "Auctions and Privatization", in *Privatization*, H. Siebert (ed.), Kiel, pp. 115-136.

[18] MYERSON, ROGER, "Optimal Auction Design", *Mathematics of Operations Research* Vol. 6, pp. 58-73.

[19] MYERSON, ROGER, AND MARK SATTERTHWAITE (1983), "Efficient Mechanisms for Bilateral Trading", *Journal of Economic Theory* Vol. 29, pp. 265-281.

[20] R DEVELOPMENT CORE TEAM (2003), *R: A Language and Environment for Statistical Computing,* R Foundation for Statistical Computing, Austria: Vienna, http://www.R-project.org.

[21] ROBERTS, KEVIN W.S. AND MARTIN L. WEITZMAN (1980), "On a General Approach to Search and Information Gathering", MIT Working Paper #263.

[22] ROCHET, JEAN-CHARLES AND PHILIPPE CHONÉ (1998), "Ironing, Sweeping, and Multidimensional Screening", *Econometrica*, Vol. 66, No.4, pp. 783-826.

[23] SAPPINGTON, DAVID E. M. (1991), "Incentives in Principal-Agent Relationships", *Journal of Economic Perspectives*, Vol. 5, No. 2, pp.45-66

[24] WEITZMAN, MARTIN L. (1979), "Optimal Search for the Best Alternative", *Econometrica*, Vol. 47(3), pp. 641-654.