ON INCENTIVES AND CONTROL IN ORGANIZATIONS

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL OF BUSINESS

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

by

Bengt Robert Holmström

December 1977

ACKNOWLEDGEMENTS

Many people and organizations have contributed to the completion of this dissertation.

Foremost, I am indebted to my principal advisor, Professor Robert B. Wilson, who has been a constant source of insight and inspiration both in courses, workshops and in his involvement with this research. His help and advise in all matters, academic and nonacademic, has been truly invaluable. Professors Charles A. Holloway and David M. Kreps, who both served on my reading committee, and Professor Joel S. Demski have all been most generous in offering me the opportunity to discuss the contents of this dissertation with them at various stages, and giving me constructive criticism. Professor Olli Lokhi (Helsinki University of Technology) has advised and aided me in my earlier years of study and also taken a keen interest in my work during my stay at Stanford.

As should be clear from reading this thesis, I have benefited much from discussions with my friends and colleagues Frøystein Gjesdal and Takao Kobayashi, to whom I owe my heartfelt thanks.

I have received financial support from Marcus Wallenberg Foundation, Suomen Kulttuurirahasto, Otto Malm Foundation, Oscar Öjlund Foundation and ASLA-Fulbright Foundation, all in Finland, and from the Graduate School of Business at Stanford.

The typing has been done by Famah Andrew with professional care.

To all these individuals and institutions I want to express my sincere gratitude.

I also want to thank my parents for their deep involvement in my studies during so many years, and my wife and son for managing to endure much neglect during the preparation of this thesis.

TABLE OF CONTENTS

Table of Contents (continued)

ON INCENTIVES AND CONTROL IN ORGANIZATIONS

Bengt Robert Holmström, Ph.D.
Stanford University, 1977

The dissertation deals with incentive problems in organizations, particularly those problems which arise due to asymmetric information among members of the organization. The incentive problem is studied in the framework of a general game-theoretic formulation. Three specific topics are discussed: delegation, coordination of information, and supply of productive inputs. These topics, though closely related, are treated essentially independently, and can be read separately.

Delegation is a decentralized decision process in which a principal lets an agent make the final decision from a restricted set of alternatives, which the principal determines. The delegation problem, from the principal's point of view, is to find the optimal set of decision alternatives to delegate to the agent. The agent is given decision-making responsibility because he has superior information about the principal's decision problem, but his freedom is generally restricted because his preferences differ from the principal's. How these two conflicting factors determine an optimal solution to the delegation problem is one of the central issues addressed. The main results on delegation demonstrate this relationship in the context of one-dimensional quantity controls. The analysis is applied to the control of centrally planned economies. It is shown that a mixed price and quantity control scheme for each agent dominates pure price schemes or quantity orders. The tightness of economic control depends on both the curvature of the benefit and cost functions, as well as the information gap between the central planner and the economic agents.

Other issues analyzed in the chapter on delegation problems include the use of general price schedules, existence of optimal delegation sets, the value of delegation, and efficiency under differential information.

Delegation is characterized by lack of communication between agents. More general decentralization mechanisms attempt to utilize simultaneously the information possessed by members of the organization. One such mechanism has been proposed by Groves. If agents have additively separable and linear preference functions, Groves' scheme will yield an efficient decision outcome. We study uniqueness properties of Groves' scheme, which complement and extend recent uniqueness results by other authors. A new technique is used to derive Groves' scheme, and this technique is subsequently employed to study the possibility of reaching efficient decisions in more complex environments.

When the supply of productive inputs cannot be observed, problems of "moral hazard" arise. Two kinds of moral hazard problems are analyzed: team production under certainty and productive services of one agent under state uncertainty.

In the case of team production it is shown that the resulting outcome will be inefficient if the outcome is shared among the input suppliers, unless a sufficiently rich measurement system is available which will discern dysfunctional behavior of agents. The problem can be avoided if an outside party is introduced, which will be used to balance the budget in a more efficient way. This amounts to a separation of ownership and labor, which provides one explanation for the existence of corporate organizations as opposed to cooperatives.

If there is uncertainty about what state of nature obtains, separation of ownership and labor will not suffice to remove inefficiencies. This

is shown in the context of a two-person principal-agent model with the agent supplying a productive input, which together with the state of nature determines the outcome. The existence of an optimal sharing rule (of the outcome) is proved and subsequently characterized. A key assumption for proving existence is that the share that goes to the agent has to be bounded. This makes the optimal sharing rule nondifferentiable. Properties of this second-best solution are studied. The analysis is extended to cases where additional observations of the agent's action and/or the state of nature are available. A necessary and sufficient condition is derived for such observations to yield Pareto-improved contracts. Verbally, the condition states that observations are of value if and only if they provide some new information about the agent's act which cannot be obtained by observing the outcome alone.

Approved for publication:

by ______________________________
For Major Department

by ______________________________
Dean of Graduate Studies

CHAPTER I

INTRODUCTION

1. Introduction

1.1 Incentive Problems in Organizations

This dissertation deals with incentive problems in organizations. There is hardly any need for motivating a study of incentive problems. The recent boom in the literature on incentives should stand as sufficient general support, and a glance at it shows that the theory is only in its initial stage with a rich set of issues uninvestigated. A study of the literature also indicates a need to unify many of the apparently diverse contributions into a more general framework, and in this way lead the research towards the central issues related to incentives. Such an attempt has been part of the objective of this dissertation, in addition to a study of some specific topics as described below.

Incentive problems arise in various organizational contexts. The incentive problems we will study, are a consequence of differential information among members of the organization. These members are viewed as self-interested economic agents who provide two kinds of services for the organization: they supply information for decision-making and inputs for production.

Because agents behave according to their own interests and since these generally differ from those of the organization, the two tasks the agents perform have to be controlled. Appropriate incentive (or control) schemes must be designed so that agents will conform as closely as possible to the organization's objectives. The difficulty lies in the asymmetry of information. In making decisions one cannot rely on the agents to tell their true information unless they are induced to do so for their own benefit. Similarly, one cannot trust the agents to provide the proper amount of inputs for production when these actions are nonobservable, unless incentive schemes are employed.[1]

In many respects the two categories -- (1) incentive problems in decision-making and (2) incentive problems in the supply of productive inputs -- are different. In the first situation asymmetric information about state uncertainty is the cause of problems; in the second situation the problem is strategic uncertainty. Since the analysis of the two problem types differ, the research on incentives has almost exclusively dealt with one or the other type in its pure form,[2] even though both types are usually present simultaneously in practice, and often intertwined in a rather complicated way.[3] This dissertation follows the same approach with a few exceptions. Chapters II and III deal essentially with pure decision-making incentives, and Chapter IV with pure production incentives.

To introduce the reader to the incentive problems we will be studying, as well as to present a unified framework for the special models we will be analyzing, we turn to a general formulation of the incentive problem.

1.2. A General Formulation of the Incentive Problem

Suppose there is a decision d to be made such that it belongs to a prespecified set of alternatives D, called the decision space. A decision process can quite generally be described as a mapping $d : M \rightarrow D$, from inputs $m \varepsilon M$ to final decisions $d(m) \varepsilon D$.

We will here be concerned with decision processes that have n participants and can be described as follows. Each participant i chooses a message m_i from his set of alternative messages M_i, called i's message space. The final decision is determined by the message n-tuple $m = (m_1, \ldots, m_n)$ via a decision function $d : M \rightarrow D$, where $M = \overset{n}{\underset{i=1}{X}} M_i$ is called the joint message space. If the decision function d depends on more than one m_i we have a decentralized decision process. The pair $N = (d,M)$ is called a decision mechanism.

We are looking at the problem of a single decision maker, called the principal, who has to choose a decision $d \varepsilon D$ facing some uncertainty. The uncertainty is described by a probability space $(Z,\mathcal{F},P)$, where $z \varepsilon Z$ is the state of nature and P is the principal's subjective probability measure on events in $\mathcal{F}$. A key feature of our problem structure is that the principal has available n agents (indexed by $i = 1, \ldots, n$), who possess some private information about the state of nature. Each agent i has observed the outcome of a random variable $\tilde{y}_i$, defined on the probability space $(Z,\mathcal{F},P)$, which provides information about $\tilde{z}$. For this reason the principal is interested in using the agents in the decision process by defining a decision mechanism $N = (d,M)$, which decentralizes the decision as described above. The agents' messages will normally correspond to part or all of their private information.

To bring out the production side of the problem explicitly, let each agent furthermore decide on a <u>productive input</u> $a_i \in A_i$.[4] A_i is called <u>i's production set</u> and $A = \overset{n}{\underset{i=1}{X}} A_i$ is called the <u>joint production set</u>. We write $a = (a_1, \ldots, a_n)$ as a generic element of A. It is assumed that the decision d is made known to the agents before they choose <u>a</u>.

What decision mechanism should the principal use? This is the main question to be addressed. The problem that the principal faces is that each agent will act in his own self-interest according to a <u>preference function</u> $F_i : D \times A \times Z \to R^1$, $i = 1, \ldots, n$, which may differ from the principal's preference function $F_0 : D \times A \times Z \to R^1$. Any particular choice of decision mechanism will result in a non-cooperative game among the agents. We will model this as a game of incomplete information (see Harsanyi [1967-1968]), so that an agent's strategy is a pair of functions $(m_i(y_i), a_i(y_i,d))$.[5] It is assumed that a Nash equilibrium provides the appropriate description of the agents' behavior, and gives the principal the basis for evaluating different decision mechanisms.

Let us now formulate the principal's problem precisely. For this we will assume that the agents know each other's preference functions F_i and the functional form of each other's private information (but not the outcomes, of course). They should also agree on the specification of the probability space $(Z,\mathcal{F},P)$. These assumptions can be defended as in Harsanyi [1967-1968], and rest on the idea that the probability space can be augmented sufficiently to incorporate all differences in information. In some cases, however, the agents will have dominant strategies, so that these assumptions need not be made. In that case agents may be completely ignorant about other agents' information structure and preferences.[6]

Given a decision mechanism $N = (d,M)$, a Nash equilibrium is defined as a set of function pairs or strategies $\{(\bar{m}_i(y_i;N),\ \bar{a}_i(y_i,d;N))\}_{i=1}^{n}$; satisfying:

(1.1) $$E[F_i(d(\bar{m}(y;N)),\ \bar{a}(y,d;N),\ z)\,|\,y_i] \geq$$

$$E[F_i(d(\bar{m}^i(y;N),\ m_i),\ \bar{a}^i(y,d;N),\ a_i,z)\,|\,y_i],$$

for every $(m_i,\ a_i) \in M_i \times A_i$, for every y_i, and

for every $i = 1, \ldots, n$.

Here, $\bar{m}(y;N) = (\bar{m}_1(y_1;N), \ldots, \bar{m}_n(y_n;N))$, $\bar{a}(y,d;N) = (\bar{a}_1(y_1,d;N), \ldots, \bar{a}_n(y_n,d;N))$. A superscript i denotes a vector with the i^{th} component deleted, e.g., $m^i = (m_1, \ldots, m_{i=1}, \ldots, m_n)$. We will sometimes write $m = (m^i,\ m_i)$.

We will assume that there exists a Nash equilibrium for each N, since this will be the case in all situations we will be studying.

In choosing among different decision mechanisms, the principal is restricted to a set of admissible decision mechanisms $\mathcal{N}$. We will have more to say about this set shortly. The principal's problem can now be stated as follows:

Decentralization Problem: Choose a decision mechanism $N \in \mathcal{N}$, such that it maximizes

$$E[F_0(d(\bar{m}(y;N)),\ \bar{a}(y,d(\bar{m}(y;N));\ N),\ z], \tag{1.2}$$

where $(\bar{m}(y;N),\ \bar{a}(y,d;N))$ is a Nash equilibrium for each N as defined in (1.1).

If there is more than one Nash equilibrium, we assume the principal will choose one of them as a basis for his optimization problem. The existence of a solution to the decentralization problem will only be discussed in connection with the special cases we will study later.

There is another way of viewing the principal's problem defined above. Let $d_0(y) = d(\bar{m}(y;N))$ and $a_0(y) = \bar{a}(y,d_0(y);\ N)$. The function part $(d_0(y),\ a_0(y))$, which is a mapping from the range of agents' observations $Y = \overset{m}{\underset{i=1}{X}} Y_i$ to D x A, is called an outcome function. The outcome function tells what decision-production pair will result when a particular decision mechanism is employed and a particular information state $y \in Y$ obtains. An outcome function is called attainable if there exists a decision mechanism $N \in \mathcal{N}$, which yields that particular outcome function at a Nash equilibrium. The set of attainable outcome functions is denoted $\mathcal{O}$. The principal's problem can then be stated equivalently as follows: find the outcome function $(d_0(y), a_0(y)) \in \mathcal{O}$, which maximizes

$$(1.3) \qquad E[F_0(d_0(y), a_0(y)), z];$$

or verbally, find the best attainable outcome function.

Since the decentralization problem is generally quite complicated, if not impossible, to solve, we will sometimes merely characterize the set of attainable outcome functions. This problem is simplified considerably by the observation that an outcome function $(d_0(y), a_0(y))$ is attainable, if and only if the decision mechanism $N = (d_0, Y)$ has a Nash equilibrium such that $\bar{m}_i(y_i) = y_i$, $i = 1, \ldots, n$, and $\bar{a}(y, d_0(y)) = a_0(y)$.[7] In other words, one can check for attainability by letting the decision part of the outcome function, $d_0(y)$, be the decision function and see if agents will tell the truth at a Nash equilibrium. It also follows that the principal may always take $M = Y$ in the decentralization problem, so that his problem reduces to a search over decision functions alone. These simple but important facts will repeatedly be used in the sequel.

Notice that our formulation does not include any explicit costs for employing different decision mechanisms. Presumably more elaborate decision mechanisms, in particular with respect to the message space M, will cost more, but have omitted such considerations.

1.3 Special Cases of the General Formulation

We now turn to some special cases of the general model described above.

Decentralization in Teams

If all preference functions F_i, $i = 0, 1, \ldots, n$, are identical, then the principal and the agents are said to form a team. The team problem has been discussed primarily in Marschak and Radner [1972] and leads to a rich theory in itself. In the team problem each agent i is in charge of a decision d_i, $i = 0, \ldots, n$, so that d can be written $d = (d_0, \ldots, d_n)$.[8] Agents receive signals y_i about the state of nature z, and subsequently may communicate some of that information to other agents according to a specific communication structure. Together the original information and the communicated information form an information system, which is simply a mapping from states of nature to states of information for each agent (i.e., n + 1 partitions of Z).

The central problem in team theory is to determine the team's best decision function d for different information systems. Based on this, alternative information systems can be compared to get an evaluation of the overall best decision mechanism.

In our framework, restriction to a particular information system can be taken into account by constraining the set of admissible decision mechanisms N. Since each agent has the same preference function, it is clear that for any optimal team decision function, the agents' message strategies constitute a Nash equilibrium.[9] It is also clear that we are not making any effective changes in the problem structure by letting the principal choose the decision functions $d_i(\cdot)$ rather than the agents. Hence, the team problem is indeed subsumed in our general formulation.

Notice that no incentive problems arise in a team, by definition. We will not study team theory in this dissertation, but it is appropriate to give a brief description of the problem, since team theory is a predecessor to recent work on incentives. Team theory for the first time recognized and modeled differential information explicitly in the context of decentralized decision-making, and this has been the basis for the development of a theory of incentives.

Delegation

Another special case arises when we restrict the decision function to have the form $d(m) = (d_1(m_1), \ldots, d_n(m_n))$. This means that each agent is assigned a separate part of the decision, which he can affect through his message. No coordination of information takes place. This corresponds to the team problem with a null communication system. Obviously, the team solution provides an upper bound for how well the principal can do with the help of his agents.

For reasons that will become evident from the discussions ahead, decentralization via a separable decision function of the form above, will be called delegation. It is a very commonly employed decentralization procedure, and will be discussed in more detail in Chapter II.

Revelation of Preferences for Public Decisions

Recently, much research has centered around the problem of designing a mechanism for public decisions, which results in Pareto optimal outcomes even when the preferences of individuals are unknown (see Groves and Loeb [1975], Groves and Ledyard [1977], Green and Laffont [1977]). The basic idea is to ask individuals for their preferences, which are parameterized by y_i; make the public decision according to a decision function $d_0(m)$, which would be Pareto optimal if the individuals told the truth (i.e., if $m_i = y_i$); and finally, induce individuals to tell the truth by a proper choice of tax functions $d_i(m)$, which are payments from individuals to the government. Since individuals are assumed by these authors to have preference functions which are linear in wealth, the transfer payments do not affect the Pareto optimality of d_0.

In our framework we would write this problem as:[10]

$$d = (d_0, d_1, \ldots, d_n),$$

$$\tilde{z} = (\tilde{y}_1, \ldots, \tilde{y}_n)$$

$$F_i(d,z) = F_i(d_0,y_i) - d_i, \quad i = 1, \ldots, n,$$

$$F_0(d,z) = \sum_{i=1}^{n} F_i(d_0,y_i).$$

Groves has shown that there exist tax functions $\{d_i\}$, such that it will be a dominant strategy for each individual to reveal his true preferences. Hence, we need not necessarily look at the problem as a game of incomplete information. However, if individuals can submit only partial information about their preferences, or if we want to eliminate one of the difficulties with Groves' scheme, called the budget-balancing problem (see footnote 10), the appropriate formulation takes the form of a game of incomplete information (see Groves [1973], d'Aspremont and Gerard-Varet [1975]).

Chapter III is devoted mostly to an analysis of revelation of preferences as described above, with a few extensions regarding the attainability of efficient outcome functions for more general preference profiles.

Moral Hazard

Moral hazard represents a pure problem of incentives for supply of productive inputs. Let us call the input to production "effort" for simplicity. There is only one agent, and both the principal and the agent have the same beliefs about the state of nature. The agent's effort $\underline{a}$ determines a monetary outcome $x(a,z)$, which also depends on the state of nature $\tilde{z}$. The problem is to determine an incentive or sharing rule $s(x)$ for the agent, which will balance gains from risk-sharing with gains from incentives to supply effort.

Formally, let $G(w)$ and $U(w,a)$ represent the principal's and the agent's utility functions over wealth and effort. The principal's decision, d_0 (which he can make himself, since the agent has no superior information about $\tilde{z}$) is the choice of the sharing rule $s(x)$. The preference function can be written:

$$F_0(d,a,z) = E[G(x(a,z) - s(x(a,z)))],$$

$$F_1(d,a,z) = E[U(s(x(a,z)), a)].$$

The principal's problem can be stated as follows. Choose $s(x)$ so that it maximizes:

$$E[G(x(a,z) - s(x(a,z)))],$$

$$\text{s.t.} \quad a = \operatorname{argmax} E[U(s(x(a,z)), a)],$$

$$E[U(s(x(a,z)), a)] \geq \bar{U}.$$

The first constraint is our Nash equilibrium condition. The second constraint is a restriction on the admissible decision $s(x)$, imposed by an equilibrium in the labor market.

In Chapter IV we will study this problem in detail and characterize the efficient sharing rules $s(x)$.

1.4 Outline of the Dissertation

The last three special structures described above represent the main topics to be analyzed in this dissertation. These topics are essentially independent. For this reason, an overview of the contents and the main results, as well as a review of the literature, have been postponed to the introduction of each chapter. Each of the Chapters II, III, and IV can also be read separately without much difficulty.

The outline of the dissertation is as follows. Chapter II deals with delegation problems, focusing particularly on the use of quantity controls in management and economic planning. Chapter III analyzes the coordination of information, mainly in the context of additive linear preference profiles as discussed above. Chapter IV studies issues of moral hazard due both to team production and to the presence of state uncertainty. Chapter V contains concluding remarks with some suggestions for future research.

Footnotes to Chapter I

[1]This may reflect an unduly restrictive view of the behavior of an agent. In particular, the possibility that moral and other reasons can make the agent conform to the organization's objectives without the use of incentive schemes is largely ignored. This is true especially when the agent's preferences are assumed to be over wealth alone.

[2]By a pure decision-making incentive problem we mean that no productive inputs are present or are assumed to cause any difficulties. By a pure production incentive problem we mean that there is no differential information about state uncertainty, so that the agents are only used for their productive inputs.

[3]Consistent with the two problems being intertwined, we find that the same incentive scheme may be used both to induce agents to tell the truth and to supply productive inputs properly.

[4]a_i could also be modeled as part of the decision d, with the agent being able to set the level of a_i independently. This formulation will be used in Chapter II.

[5]a_i will depend on d, since we assumed the agent knows d at the time he determines a_i.

[6]Usually the game is no longer called a game of incomplete information when dominant strategies give the solution.

[7]To see this, let $(\bar{m}(y;N), \bar{a}(y,d;N))$ be a Nash equilibrium with a decision mechanism $N = (d,M)$. Take $N_0 = (d_0, Y)$, with $d_0 = d_0 \circ \bar{m}$, and this will have a Nash equilibrium with $m_i(y_i) = y_i$, $\bar{a}(y,d_0(y)) = a_0(y)$. These Nash equilibria are essentially identical, since they result in the same outcome function. From a practical point of view, they are, of course, different in terms of communication costs. Often the decision function is such that not all of the private information of agents is needed.

[8]Since everybody has the same objective function, we can take the a_i's to be part of the decisions d_i.

[9]With concave preferences the reverse is true. A Nash equilibrium gives the optimal team decision function. This is one of the main theorems in team theory (see Marschak and Radner [1972]).

[10]The tax functions are not included in F_0, and consequently only d_0 will be efficient, unless the taxes happen to sum to zero. The question whether there exist taxes which sum up to zero and achieve an efficient outcome function d_0 is known as the budget-balancing problem. Generally it cannot be solved, as we will see in Chapter III.

CHAPTER II

DELEGATION

2.1 Introduction

In this chapter we will look at a special but important type of decentralized process, called delegation. Delegation is formally defined by the characteristic that the decision function separates so that each agent affects only his own decision and no coordination of information takes place. As we will see shortly, this is equivalent to a process in which agents are given a set of alternatives from which they can choose according to their own preferences; in other words, the agents are delegated the final decision, possibly subject to some constraints imposed by the principal.

Recently, many authors have discovered and explained the delegation process in various economic contexts; sometimes so well disguised that it has not been immediately clear that the same basic phenomenon has been studied. One of the earliest, if not the earliest, paper on this subject is by Mirrlees [1971] and deals with optimal taxation. The government delegates partly the decision of how much taxes individuals are to pay, by tying the taxes to the amount individuals earn. The reason is that the government, according to Mirrlees' welfare criterion, ideally would tax individuals based on their abilities. But the government does not know these abilities, and so by delegating the decision it can improve on a centralized solution.

More familiar is the interpretation that the government will screen individuals by levying an income tax. Individuals with different abilities will choose different levels of income and thus differentiate themselves according to ability. This self-screening process can also be called signaling. Individuals will reveal their abilities via their choice behavior.

The research into screening and signaling phenomena (which are cases of delegation) gained substantial impetus with the path-breaking dissertation of Spence [1973] on signaling in labor markets and with independent work on screening in insurance markets by Rothschild and Stiglitz [1976]. These studies differ from Mirrlees' taxation problem in that they analyze informational equilibria in markets, and how market performance is affected by asymmetric information.[1] Both Spence and Stiglitz (together with several others) identify a large number of instances in which signaling or screening play a crucial role, for example in labor markets, insurance markets and capital markets. This should come as no surprise, since informational asymmetry is the rule rather than the exception in the real world.

The results of these studies into the effects of asymmetric information on market operation have been surprising in many respects. One finds that ordinary price equilibria are no longer viable and are replaced by prices supplemented with quantity rationing, or more generally by nonlinear price schedules. The market outcomes are inefficient with respect to the standard of perfect information, because transmission of information becomes costly in contrast to costless information transmission (about preferences) in classical markets.[2] Multiple equilibria may exist, but what is more puzzling and controversial, there may be no equilibrium. This, of course, suggests that the models used must be incomplete in some respect, since nonexistence of equilibrium is hardly acceptable. But aside from this issue, the models have greatly enriched our views of how delegation operates in the economy as a means of transmitting information and overcoming informational gaps. They have explained many of the peculiar institutions formed in markets with incomplete information.

Though signaling and screening are cases of delegation, we will not be studying the market models referred to above. Our interest here lies in the use of delegation in centrally planned economies and in management of firms. But similar processes are at work in both types of economies.

Since information is not coordinated when delegation is used, much insight can be gained by studying a simple two-person principal-agent relationship. This we will do for a large part of the chapter. Our objective is to understand how delegation improves on decision-making and what determines the forms it takes.

The outline of the chapter is as follows. We begin by formulating the delegation problem from the principal's point of view, and prove some general theorems about the existence of an optimal solution. Several examples of delegation are given to illustrate the formulation. The core of the chapter is an analysis of a particularly simple form of delegation; namely quantity controls. Because of their simplicity and minimal informational requirements, they play an important role in management and economic control. We also look at more general price schedules to see to what extent economic agents can be controlled, and apply the results to a problem of economic regulation. Finally, we analyze the efficiency properties of the solution to the delegation problem, and discuss in general the difficulties of finding a viable notion of efficiency under differential information. The chapter concludes with some remarks about extensions to dynamic problems and issues of observability.

2.2 The Two-Person, Principal-Agent Case

2.2.1 Formulation of the Delegation Problem

The delegation problem can be written somewhat differently than our general problem formulation in Chapter I, in order to emphasize the specific nature of delegation and to get a more convenient model. We will give the delegation formulation first and later show that it is, in fact, equivalent to the general formulation given earlier. To this end, define the following concepts (some of them known from before):

D — the decision set; a compact subset of a complete metric space. A generic element of D is denoted by d and called a decision or alternatively an action.[3]

$\mathcal{C}$ — the set of admissible controls; a closed set of nonempty bounded closed subsets of D, where closedness of $\mathcal{C}$ is w.r,t., the Hausdorff-metric. As is seen in Appendix 2A, $\mathcal{C}$ is then compact in this metric. A generic element of $\mathcal{C}$ is denoted C and called a control (or control set).

$(Z,\mathcal{F},P)$ — a probability triple, describing the state of nature $\tilde{z}$, with Z a topological space.

$F^P(d,z)$
$(F^A(d,z))$ — the principal's (the agent's) preference function, which is assumed jointly continuous in both arguments and uniformly bounded.

$\tilde{y}$ — the agent's private information; a random variable on $(Z,\mathcal{F})$ (possibly vector valued). The range of $\tilde{y}$ is denoted Y, and assumed to be a subset of a topological space.

We will generally use the symbol P for probability measures. The argument should reveal which particular probability measure is relevant. For instance P(y) is the induced probability measure of y. $P(z|y)$ is the conditional probability measure of z given $\tilde{y} = y$, and so forth.

It is assumed that the principal and the agent agree on the structure of $(Z,\mathcal{F},P)$. However, the agent has superior information about $\tilde{z}$, in that he has observed the outcome of $\tilde{y}$.[4]

With these definitions the delegation process can be described as follows:

1. The principal chooses a control $C \epsilon \mathcal{C}$;

2. Then, knowing C and y, the agent chooses an action $d \epsilon C$.

Subsequently payoffs $F^P(d,z)$ and $F^A(d,z)$ are realized.

Schematically this looks like:

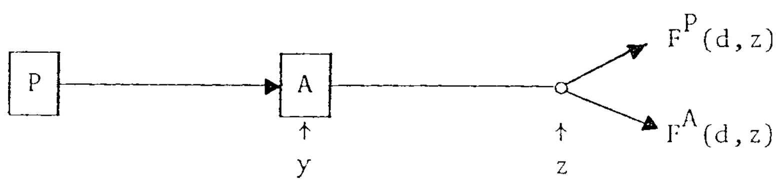

The principal delegates the final decision to the agent by choosing a control set C, from which the agent is allowed to pick any action he wants.

What C should the principal choose? To formulate this problem, we examine the agent's reaction to a particular control C. Given the outcome of $\tilde{y}$, the agent will determine d by solving[5]

$$(2.1) \qquad \max_{d \epsilon C} \; E[F^A(d,z)\,|\,y].$$

This maximization is well defined since C is a closed subset of a compact set D and F^A is continuous. The solution may be non-unique, in which case we have to specify how the agent makes his final choice. Let $d(y,C)$ be the solution set to (2.1), called the <u>solution</u>

correspondence. For a fixed C it is a correspondence in y. Any particular representation of d(y,C), i.e., any function of y which for each y takes a value in the solution set d(y,C), will be called a controlled response function and denoted $d(y|C)$. If C = D, we will write $d(y|D) = d(y)$ and call it the uncontrolled response function (when no confusion will arise we omit the qualifiers "controlled" or "uncontrolled"). Notice that even if the principal does not observe y, a response function can still be defined for him corresponding to (2.1), and this fact will frequently be used.

Two particular response functions one could think of as a basis for the principal's choice of C, are:

(2.2) $$d_{min}(y|C) = \text{Argmin } \{E[F^P(d,z)|y] \mid d\varepsilon d(y,C)\}, \ \forall y,$$

or

$$d_{max}(y|c) = \text{Argmax } \{E[F^P(d,z)|y] \mid d\varepsilon d(y,C)\}, \ \forall y.$$

In (2.2) the principal is pessimistic about the agent's choice in case of multiple maxima in (2.1), and calculates according to the worst possibility. In (2.3) the principal is as optimistic as possible.

We will always assume, unless otherwise stated, that the agent's as well as the principal's uncontrolled response function is continuous in the signal y. This would be implied by the assumption that $E(F^A(d,z)|y)$ $(E(F^P(d,z)|y))$ is continuous in d and y and strictly concave in d for each y, and D is convex.

The principal's problem can now be stated as follows:

Delegation Problem (DP)

Find $C^* \in \mathcal{C}$, which maximizes

$$E[F^P(d(y|C),z)], \tag{2.4}$$

where $d(y|C)$ is a controlled response function of the agent.

For notational convenience we will write:

$$f^i(d,y) = E[F^i(d,z)|y], \qquad i = A, P.$$

Note that $f^i(d,y)$ is continuous in both arguments and uniformly bounded. Also, if F^i is concave in d for each y, so is f^i. Occasionally, we will work with loss functions rather than preference functions. These are defined as the negatives of the preference functions and denoted by L^i respectively ℓ^i.

2.2.2 Existence of a Solution

It is not a priori clear that a solution C^* exists; in fact, it does not in general, as the following counter example shows.

Example 2.1: Let the agent respond according to (2.2). Take $\tilde{y} = \tilde{z}$; $\tilde{y}$ can take on two values, 0 or 2, with equal probabilities. $D = [0, 10]$, $\mathcal{C} = \{\{d,z\}|a = 0, b\varepsilon D\}$. The only control parameter is thus $b \varepsilon D$. Let $F^A(d,z) = -(d - z - 1)^2$, and $F^P(d,z) = -(d - z)^2$. The agent's and the principal's uncontrolled response functions are depicted below.

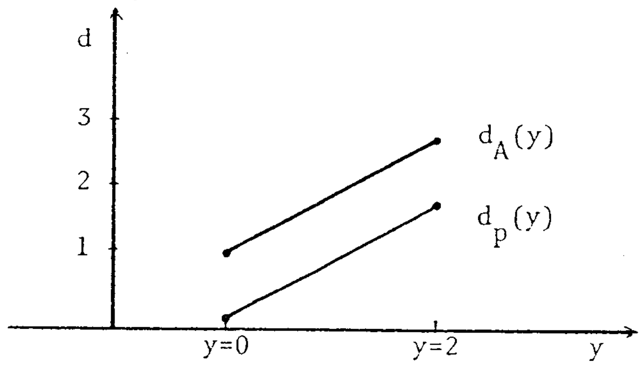

The supremum of the control problem is 0. It can be approximated arbitrarily closely by taking $b = 2 + \varepsilon$. With the response function d_{min}, however, the supremum cannot be attained, since $b = 2$ will result in the agent picking $d = 2$ both if $y = 0$ and if $y = 2$. The problem is that as $b \to 2$ from above, the agent's response function will make a jump at $b = 2$. Moreover, this discontinuity occurs on a set of probability measure 1/2, and so the discontinuity is carried over to the principal's objective function. □

The example indicates two ways of amending the problem. One way is to assume that the agent responds according to the optimistic function in (2.3). In the example this would have implied that the agent would have chosen $d = 0$ for $y = 0$ and $d = 2$ for $y = 2$ given the control $\{0,2\}$, and this would have attained the supremum value. The other way is to assume that discontinuities occur only on sets of measure zero. For instance, if the principal would have been uncertain about the agent's preference function, described by an added term of uncertainty in F^A, say a uniformly distributed noise term, then it is readily seen that there would have existed an optimal control in the example.

In the appendix the following two theorems are proved. They show that both ways of amending the problem work.

Theorem 3.1: Assume that Prob $\{d(y,C)$ is not a singleton$\} = 0$ for every $C \varepsilon \mathcal{C}$. Then there exists an optimal solution to the delegation problem (2.4), regardless of which particular response function the agent uses.

Theorem 3.2: For the response function $d_{max}(y|C)$, there exists an optimal solution to the delegation problem.[6]

The assumption in Theorem 3.1 is generally satisfied if the principal's model about the agent's behavior is imperfect (see footnote 4).

2.2.3 Connection to the General Decentralization Formulation

Before discussing some examples of delegation, let us compare the delegation formulation (2.4) with the general decentralization problem (1.2) and check that the former is a special case of the latter. In the decentralization problem the principal is looking for a decision mechanism $N = (d,M)$. For any particular choice of N, write

$$(2.5) \qquad C = \{d(m)\varepsilon D | m\varepsilon M\} \subset D.$$

Since the agent can choose any message $m\varepsilon M$ he wants to according to the decentralization formulation, he has in effect the freedom to pick any decision in C defined by (2.5). Hence, a decision on N amounts to specifying a control set C, which is a subset of D. The ensuing game of incomplete information, defined in (1.1), degenerates to the agent's maximizing behavior (2.1), and the set of admissible decision mechanisms in the decentralization problem is represented by the set of admissible controls in the delegation formulation.

On the other hand, it is immediate that a delegation problem can be formulated as a decentralization problem by taking the decision itself as a message. That is, for an arbitrary control set C, let $M = C$ and take $d(m) = m$ as the decision function.

This discussion shows that in the context of a principal-agent game, decentralization through communication and through delegation are equivalent.

2.2.4 Some Examples of Delegation

As formulated, the delegation problem is quite general. We have written it as compactly as possible to reveal the basic structure of delegation. It also allowed us to prove quite general theorems on existence of a solution. On the other hand, it may be somewhat confusing to work with preference functions directly defined over decisions and states of nature, since in general these are measures derived from some underlying problem structure, combined with preference measures independent of the problem itself (like a utility function over wealth alone). To clarify the situation as well as to show the applicability of our formulation let us look at some examples.

Example 2.2: A Production Problem

In this example, the agent is a divisional manager and the principal is a representative for the headquarters in a firm. The problem is to decide on the production level d for one of the division's products. The price of the product is known and for simplicity assumed equal to 1. The cost function is uncertain and defined by

$$C(d,\tilde{z}) = d^2 - 2\tilde{z}\cdot d,$$

where $\tilde{z}$ is a random variable.[7] The agent has superior information about $\tilde{z}$ in that $\tilde{z} = \tilde{y} + \tilde{x}$, and the agent knows the outcome of $\tilde{y}$. The principal does not know $\tilde{y}$ and neither one knows $\tilde{x}$.

The principal would like the agent to make the decision, since the agent is better informed, but the problem is that the agent's and the principal's preferences differ. Both have a multi-attribute utility function of the form b • sales + profit, but with differing weights b on sales. In particular, we will assume that the division puts more weight on sales than the center, since it looks narrowly at its own growth objectives, whereas the center is concerned with the firm's overall performance (this assumption plays no crucial role; only the fact that preferences differ is important). We can then write the preference functions as

$$(2.6) \qquad \begin{aligned} F^P(d,z) &= \bar{b}_p \cdot d - d^2 + 2z \cdot d, \\ F^A(d,z) &= \bar{b}_A \cdot d - d^2 + 2z \cdot d. \end{aligned} \qquad \bar{b}_A > \bar{b}_p ,$$

The principal's problem is to decide on how much freedom the agent should be given in choosing d. We will return to the solution of this problem in section 2.3.2, which will illustrate the basic trade-off between differences in information and in preferences when delegating. □

Example 2.3: An Insurance Model

As we mentioned in the introduction of this chapter, screening and signaling models can be viewed as particular forms of delegation. Though these models deal with more than two persons, they normally involve only bilateral trade, which can be put in the framework of our principal-agent paradigm. In order to be specific, we will look at the case of an insurance market and mention in passing how our framework applies to other models of screening.

The insurance company is the principal and the insured the agent. In the simplest case, the agent faces a risk of having an accident, described by a random variable $\tilde{x}$ (= 0 if no accident, = 1 if accident). The agent knows his probability of having an accident, denoted by y, whereas the insurance company is uncertain about y. We can then take $\tilde{z} = (\tilde{x},\tilde{y})$. The decision is what insurance policy d should be offered to the agent. d can be written as d = (a,t), where a is the level of insurance (payment in case of an accident, net of premium) and t is the premium. The company is assumed to be risk-neutral, and has a preference function

$$F^P(d,z) = t \quad , \text{ if } \tilde{x} = 0$$
$$= -a \quad , \text{ if } \tilde{x} = 1.$$

The agent is risk-averse, with a utility function U and consequently a preference function

$$F^A(d,z) = U(-t) \quad , \text{ if } \tilde{x} = 0$$
$$= U(a) \quad , \text{ if } \tilde{x} = 1.$$

Since the company does not know $\tilde{y}$, it turns out that it is beneficial to delegate the choice of an insurance policy to the agent by offering him a set C of insurance policies, rather than a single one. The agent prefers, for any fixed level of insurance a, smaller premiums to larger, so the company only needs to consider delegation of sets of the form $C = \{(a, s(a))\}$. This can be seen from the picture below.

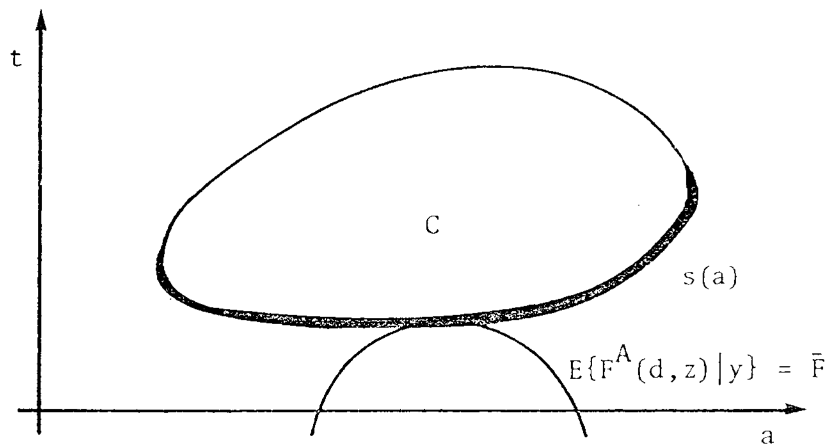

$s(a)$ is the price the agent has to pay for the level $\underline{a}$ of insurance which he chooses from C. We see that the company's optimal pricing problem is one of optimal delegation.

The labor market model treated by Spence [1973] and the taxation problem of Mirrlees [1971] are very similar to the insurance model above. In Spence's model the principal is a firm, and the agent is a worker. $\tilde{y}$ is the "ability" of the worker, which only he knows; $\underline{a}$ is the level of education he chooses to purchase. The company will screen individuals by offering a nonlinear wage schedule $s(a)$ based on education.

In the taxation problem the government is the principal, and a citizen is the agent. $\tilde{y}$ stands again for ability, $\underline{a}$ for income, and the problem is to design an optimal tax function $s(a)$. The objective for the government is some welfare function, which it tries to maximize subject to the constraint of breaking even.

It should be noted that the market models under incomplete information have a special feature, which does not appear in a centrally planned economy. The agent always has available a set of contracts offered by the other firms, and this means that when one company plans for its optimal set of offers, it has to consider the other companies' current offers. Our formulation should be slightly changed to take this into account. Simply redefine the agent's response function so that he chooses among $C \cup C_0$, where C is the principal's set of offers and C_0 is the set of offers from the other companies. If the agent picks $d \in C_0$, this implies that no trade occurs between the principal and the agent. The existence of C_0 is a source of problems for the market models, since it leads to nonexistence of equilibria (see Riley [1976]).

We will be discussing screening models in more detail in Section 2.4.

Example 2.4: Management by Participation

In Weitzman [1976a] we find a description of a new incentive structure that has been introduced in the Soviet Union for control of production units. Originally, the center in the planned economy announced fixed targets for each production unit, with bonuses for overfulfillment and penalties for underfulfillment. In 1972 the system was changed so that the production units could make changes in the targets at a certain cost. Weitzman explains this as an attempt to solve the dynamic incentive problem more efficiently. In our interpretation the reason for the change is that the optimal production target is a function of information that only the production unit knows. In particular, it is a function of the current output potential and the disutility of effort, and as an optimal way of eliciting effort the targets are partly allowed to be set by the units themselves. This is what goes on in most western companies and can be called management by participation.

Briefly described, the Soviet scheme looks as follows. The final reward to a unit is given by:

$$\begin{aligned} (2.7) \qquad R(t,g,x) &= B + \beta(g-t) + \alpha(x-g) \quad , \text{ if } x \geq g, \\ &= B + \beta(g-t) - \gamma(g-x) \quad , \text{ if } x < g, \end{aligned}$$

where

x = final production,

g = the goal, set by the production unit,

t = the original target, set by the center,

B = fixed bonus.

The parameters satisfy $0 < \alpha < \beta < \gamma$. From (2.7) we can see that the original target t can be changed to g at a price $\beta\cdot$ per unit. Overfulfillment is rewarded at a rate α, and underfulfillment is penalized at a rate γ. If the production function is $x = x(e,\tilde{z})$ where e is effort (or other inputs) and $\tilde{z}$ is uncertainty, and $u(w,e)$ is the agent's utility function over wealth and effort, we get a delegation problem with

$$(2.8) \qquad F^A(g,e,z) = U(R(t,g,x(e,z)),e),$$

$$F^P(g,e,z) = x(e,z).$$

(The center may, of course, have some other objective than this.) The decision is d = (d,e), of which only g is controllable by the center. This problem can also be seen as an example of screening, but with a linear trade-off function between the primary decision variable g and the compensation R. It will become clear later that with (2.7) and (2.8) as our analytic model, it is rational to let the production units partly decide on their own goal.[8] □

Our final example deals with a generalization of one of the earliest incentive models studied.

Example 2.5: Risk Incentives (see Wilson [1968], Ross [1973]).

The agent's utility function is $U^A(.)$ and the principal's $U^P(.)$; both are defined over wealth alone. The principal faces a decision a with an uncertain monetary outcome $x(a,\tilde{z})$. This decision is delegated to the agent, since he has superior information about $\tilde{z}$. To provide the agent with proper incentives he is given a share in the outcome, written s(x,a), which may depend on the decision a directly only if it is observable. The problem is then:

$$(2.9) \qquad \max_{s(x,a)} \int U^P(x(a(y),z) - s(x(a(y),z),\ a(y)))dP(z,y),$$

$$\text{s.t.} \qquad a(y) = \underset{a}{\operatorname{argmax}} \int U^A(s(x(a,z),a))dP(z,y)\ \forall y,$$

$$\int U^A(s(x(a(y),z),a(y))dP(z,y) \geq \bar{U}.$$

Previously, in the work of Wilson [1968] and Ross [1973], differential information has only been considered implicitly. The reason is that their interest has been focused on choosing s so that the agent's and the principal's induced utility functions become identical. In that case the distribution of $\tilde{y}$ does not matter. Moreover s need not have a as an explicit argument. If there is no asymmetry in information, then the principal himself can specify the best sharing rule s (and the best act a to be taken, if a is observable); i.e., pick the best decision d = (s,a), where s is a function. With differential information it pays in general to decentralize both decisions.

The derived preference functions F^A and F^P are clear from the expressions in (2.9). □

From this list of examples it should be clear that the delegation problem is quite general. It could have been written in a more extensive form earlier, but we chose a compact formulation for ease of exposition. From the examples we see that the decision can generally be written as $d = (a,e,w(\tilde{x}))$, where a is the primary decision to be made (e.g., production level, amount of insurance), e is a nonobservable act of the agent which cannot be controlled (e.g., effort), and w is compensation of some form (e.g., tax payments, insurance premiums), which may depend on a random variable $\tilde{x}$, which is observable ex post (as in Example 2.5).

2.2.5 The Admissible Set of Controls

A number of considerations influence the specification of the admissible control set C. Observability is one. The principal cannot restrict the agent to a set C unless he can observe whether the agent takes a decision in C or not; and naturally the agent has also to be able to observe the same. Secondly, market forces may impose restrictions on C (as was seen in Example 2.3). The principal-agent model is a partial-equilibrium model and one would normally suppose the agent (or the principal) will have requirements on a minimum expected utility level. Thirdly, one might want to restrict C to contain only certain simple forms of controls, due to costs of using other and more complicated forms or due to the fact that the delegation problem is too hard to solve in general. This is, for instance, the approach taken in the theory of teams, in the Soviet incentive model (Weitzman [1976a]), and in Weitzman [1974], where only prices and quantities are compared.

Likewise, we will look at some special control forms and problem structures. We start by studying the one-dimensional control problem.

2.3 Quantity Controls

In the simplest form of delegation the control parameter is one-dimensional and the controls are subsets of $\mathbb{R}^1$. Such controls we call quantity controls. Quantity controls are widely used in practice. We see them in the management of firms; for instance, when a division is allowed to make investments up to a certain limit on its own, but beyond it only with the center's approval; or when a division's production is required to lie within certain limits for coordination with other divisions. Other examples would be a person who lets a broker manage a limited amount of his funds, or a bank which regulates the amount a person can borrow. Quantity controls are also implicit in the market economy; whenever the market is in disequilibrium the long side is rationed by the short side.

Various reasons can be given for the use of quantity controls. Our main interest was in explaining them as tools for utilizing the agents' information and expertise in organizations, as well as for coordination of decision making. Quantity controls lend themselves well to studying how the two central components of decentralization -- the difference in information and the difference in preferences -- affect the optimal delegation decision, and this is a point we will stress throughout.

We mentioned already that delegation economizes on communication costs. The agents need not communicate with the principal in order to know what they should do. However, the principal should still be able to observe the agents' decision d to make sure it lies in the control set C, though with quantity controls such monitoring may well be less costly than when more elaborate controls are used.[9]

Though we will deal predominantly with differential information as described in the basic model of delegation, we start with showing that quantity controls have a natural place as complements to prices in an imperfect market. We will then study extensively the case where preferences are or can adequately be approximated by quadratic functions, and give interpretations of the results in the context of management control. The same model will be applied to analyze an extension of Weitzman's well-known paper on comparison of prices vs. quantities [1974]. We also ask to what extent the intuitive results from the special quadratic cases generalize, and study the basic question: when does it pay to decentralize via quantity controls. Finally, we exhibit the role of delegation as a motivator for the agents to collect more information (thus giving an analytic interpretation of one often stated advantage of decentralization), and as a tool for receiving information about the agents' characteristics.

2.3.1 Quantity Controls as Substitutes for Market Imperfections

Kurz [1976] has studied rationing in credit markets as an example of how quantity controls can be used to compensate for a market imperfection. We will be brief and only discuss the main point in the argument.

Consider a firm which wants to borrow money from a bank for a given project. The project yields constant stochastic returns to scale, and is characterized by the state variable $\tilde{z}$ = 1 + rate of return; $z \varepsilon [0,\infty)$, $G(z)$ is the distribution function and $g(z)$ the density of $\tilde{z}$. The firm can put up K dollars for the project, which also represents its maximum liability for repaying the loan. If the state falls below a critical value $\bar{z}$, which depends on the amount of the loan d, the interest rate r, as well as K, the firm will go bankrupt, and this possibility causes the market imperfection, since there is no market for bankruptcy insurance. The bank is assumed risk neutral and the firm risk-averse with utility function U; $U' > 0$, $U'' < 0$. Denote the payoff functions of the bank and the firm F_B, F_F, respectively. We have

$$F_B = \begin{cases} r \cdot d & , z \geq \bar{z}, \\ (z - 1)d + z \cdot K & , z < \bar{z}. \end{cases}$$

$$F_F = \begin{cases} z(d + K) - (1 + r) \cdot d - K & , z \geq \bar{z}, \\ -K & , z < \bar{z}, \end{cases}$$

where

$$\bar{z} = (1 + r) \frac{d}{d + K} .$$

If the bank could only charge a certain interest rate r, but not control the amount the firm borrows, the firm would demand a loan of size d, which solves the program:

$$\max_d E_F(d,r) \equiv U(-K) \cdot P(z < \bar{z}) + \int_{\bar{z}}^{\infty} U[z(d+K) - (1+r)d - K]dG(z).$$

The first-order condition is:

$$\int_{\bar{z}}^{\infty} U'[z(d+K) - (1+r)d - K](z - (1+r))dG(z) = 0.$$

This equation describes the demand function d(r), which decreases in r.[10]

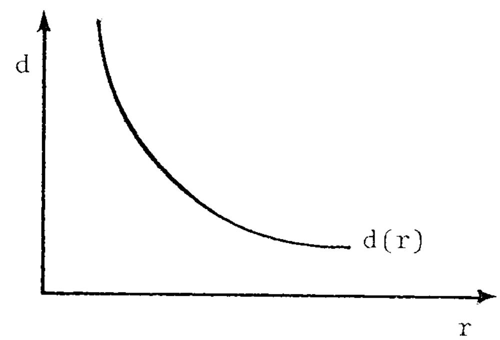

Each point in the (d,r) - space pictured above corresponds to an expected utility pair $(E_B(d,r), E_F(d,r))$ of the bank and the firm, respectively. With pure interest contracts only points along the demand curve can be reached, whereas with quantity controls, any point in the (d,r) - space is available. The purpose is to show that almost every contract (r, d(r)) can be strictly Pareto-dominated by a contract of the more general form (r,d).

For this to be true, it suffices to show that the gradients of the expected utility functions E_B and E_F are linearly independent along the demand curve. And this is the case if $\partial E_B/\partial d \neq 0$ along $d(r)$, since by definition $\partial E_F/\partial d = 0$ along $d(r)$. We compute:

$$\left.\frac{\partial E_B}{\partial d}\right|_{(r,d(r))} = r \cdot P(z \geq \bar{z}) + \int_0^{\bar{z}} (z - 1)dG(z).$$

The first term is positive whereas the second term may be negative so, indeed, this expression can equal zero. However, this happens on a negligible set in the appropriate space of probability distributions $G(z)$ and utility functions U. For this reason, almost every point along the demand curve $d(r)$ can be dominated by a contract which includes quantity constraints. This seems a convincing reason for the presence of quantity controls in the loan market. Moreover, it is readily checked that if we are at an equilibrium in a market which uses only prices (i.e., interest rates), then the zero profit condition for the bank will imply that

$$\frac{\partial E_B}{\partial d} = \frac{-1}{d} \int_0^{\bar{z}} z \cdot K \cdot dG(z) < 0.$$

Hence, a pure price equilibrium can always be improved upon by additional rationing.

Note that there is no differential information about the return of the project. If the firm would have superior information about $\tilde{z}$, this would give a further reason for using quantity controls, in particular interval controls, as will become clear shortly.

2.3.2 Quantity Controls in Management

We begin by analyzing extensively the production problem formulated in Example 2.2.

Example 2.2 (continued):

The preference functions that were defined for the center and the division by (2.6) can be rewritten as the following loss functions:

$$L_P(d,z) = (d - b_P - z)^2,$$

$$b_A > b_P.$$

$$L_A(d,z) = (d - b_A - z)^2,$$

by completing the square (constants will not affect the decisions), and letting $b_i = \frac{1}{2}\bar{b}_i$, $i = A,P$. Recall that $\tilde{z} = \tilde{x} + \tilde{y}$ and the division knows the outcome of $\tilde{y}$. We will assume that $\tilde{x}$ and $\tilde{y}$ are independent random variables. $\tilde{y}$ has a uniform distribution over $(-\delta,\delta)$ and $\tilde{x}$ has an arbitrary distribution with finite mean m and variance s^2.

The uncontrolled response function of the agent and the principal are:

$$d_A(y) = b_A + m + y,$$

$$d_P(y) = b_P + m + y, \quad \text{(if the principal knew y).}$$

Since $b_A > b_P$, the agent is constantly biased towards choosing a too-high production level d. In the figure below we have drawn the uncontrolled response function $d_A(y|C)$ with $C = (-\infty, d_U)$.

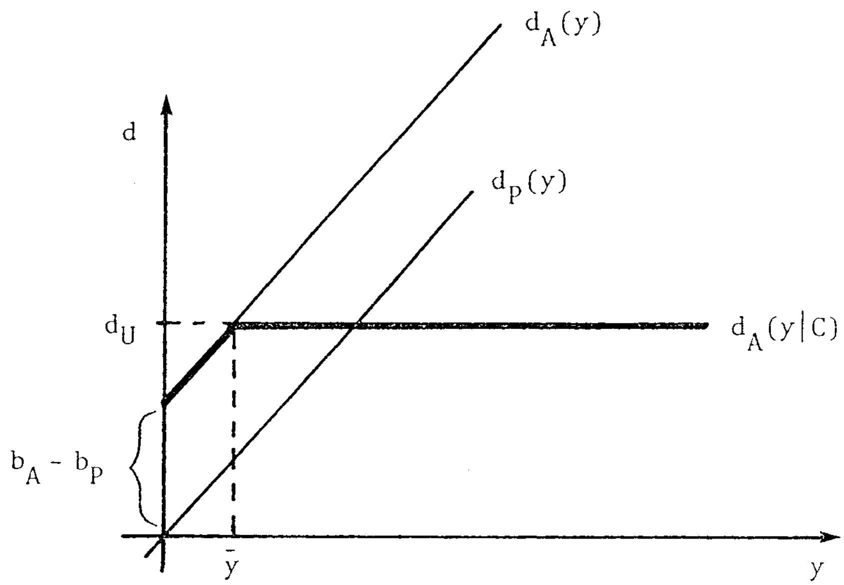

We will first calculate the optimal interval control. From the fact that the agent is constantly biased towards higher d values, it follows that an optimal interval control takes the form $(-\infty, d_U]$. We only have to find the optimal value for d_U. Since $d_A(y)$ is increasing, this is equivalent to finding a critical value y_U of y, such that for $y \leq y_U$ the agent can exercise his preferred act, and for $y > y_U$ he chooses d_U, where $d_U = d_A(y_U)$. It will be more convenient to find the optimal value for y_U.

For a particular y_U, the principal's expected loss is:

$$E_P(y_U) = \int_{-\delta}^{y_U} \int_{-\infty}^{\infty} (d_A(y) - d_P(y))^2 dP(x)dP(y),$$

$$+ \int_{y_U}^{\delta} \int_{-\infty}^{\infty} (d_U - d_P(y))^2 dP(x)dP(y)$$

$$= (b_A - b_P)^2 + s^2 + \int_{y_U}^{\delta} [2(b_A - b_P)(y_U - y) + (y_U - y)^2]dP(y).$$

Since the principal's preference function is quadratic, only the mean and variance of $\tilde{x}$ are relevant. It is a straightforward exercise to calculate the optimal y_U^* by differentiating $E_P(y_U)$. Using $d_U^* = d_A(y_U^*)$, the result is:

$$\text{(2.10)} \qquad d_U^* = b_P + m + \max(0, \delta - (b_A - b_P)).$$

The best centralized act by the principal is $d^* = b_P + m$, which is always contained in $(-\infty, d_U^*)$. δ is a measure of the information gap between the principal and the agent, and $(b_A - b_P)$ measures the difference in preferences. Hence, the result is intuitively appealing: the agent is given more freedom with an increased information gap or with closer preferences. This is what we would expect as an illustration of the fundamental tradeoff between differences in objectives and in information when decentralizing. It is natural to ask how generally this conclusion is valid, and this will be discussed in the later parts of the section.

We will see that generally the optimal quantity control does not take the form of an interval. In this particular case, however, it is easy to see that the interval $(-\infty, d_U^*)$ is the overall optimal quantity control. Make the contrapositive assumption that an optimal control has a gap (d_1, d_2) somewhere, i.e., that the agent is not allowed to use values in this interval but can use values above or below. Compare such a control to one where the gap is "filled in." Since the agent's response function is increasing we can make this comparison over the corresponding gap (y_1, y_2), where $d_1 = d_A(y_1)$, $d_2 = d_A(y_2)$. A direct evaluation of the principal's losses with the two controls shows that the control with the gap is inferior by the amount $\frac{1}{12} \cdot \frac{1}{2\delta} \cdot (y_2 - y_1)^3 > 0$.

It should be noticed that one has to be somewhat careful with the interpretation of a change in δ. δ is a measure of the information gap, but what does it mean that δ increases? It cannot be interpreted so that the agent collects some further information about $\tilde{z}$, since with a change in δ, the principal's beliefs about $\tilde{z}$ will change, which is inconsistent when his information actually stays the same.

The proper interpretation is that we are looking at two repeated situations. In the second one the principal is less informed than in the first one, whereas the agent has the same information as before. An example would be the following situation. The principal knew yesterday a relevant random variable (e.g., a stock price), but today he only has a uniform distribution on it, whereas the agent has observed the actual outcome both days. Our example indicates that in such a situation the principal will give the agent more freedom today than yesterday.

When we want to analyze a situation in which the agent gets more informed, while the principal's information stays the same, we can look at normal distributions. Think of $\tilde{z}$ as the sum of a large number of independent random variables. The agent observes part of them, and this part sums up to $\tilde{y}$ while the remainder is $\tilde{x}$. By observing in addition a few of the variables that go into $\tilde{x}$, the agent gets more informed. Assume the number of random variables in $\tilde{x}$ and $\tilde{y}$ are so large that they can be considered normally distributed. Let $E_x = m$, $E_y = 0$ (w.l.o.g), Var $(x) = s_x^2$, Var $(y) = s_y^2$. Then $\tilde{z} \sim N(m,s^2)$, $s^2 = s_x^2 + s_y^2$. The agent's acquisition of new information (while the principal's knowledge stays the same), can adequately be described as an increase in s_y^2 and a similar decrease in s_x^2, with s^2 staying constant. On the other hand, if the principal's information changes but the agent's does not, then s_y^2 changes with s_x^2 constant. This has the interpretation we gave above for the uniform distribution. With quadratic loss functions only the change in s_y^2 matters for determining the optimal control, and so the analysis is the same regardless of which interpretation we want to use.

Let us now verify that the qualitative conclusions we made in the uniform case are also true for the information structure described above with $\tilde{z} = \tilde{x} + \tilde{y}$, all normally distributed.[11] We will work with a more general quadratic preference structure for later applications.

Let the loss functions for the agent and the principal be:

(2.11)
$$L_P(d,z) = (d - (a_P \cdot z + b_P))^2,$$
$$L_A(d,z) = (d - (a_A \cdot z + b_A))^2.$$

Conditional on the observation y, the expected losses are:

(2.12)
$$\ell_P(d,y) = (d - (a_P(y+m) + b_P))^2 + a_P^2 \cdot s_x^2,$$
$$\ell_A(d,y) = (d - (a_A(y+m) + b_A))^2 + a_A^2 \cdot s_x^2.$$

Consequently, the uncontrolled response functions are:

(2.13)
$$d_P(y) = a_P(y+m) + b_P,$$
$$d_A(y) = a_A(y+m) + b_A.$$

We will assume that $a_A \geq a_P > 0$ and $b_A \geq b_P$. The other cases can be analyzed in a similar manner.

Let (d_L, d_U) be a control interval. The expected loss for the principal when he uses this control is:

(2.14)
$$E_P(d_L,d_U) = \int_{-\infty}^{y_L} (d_L - d_P(y))^2 dG(y) + \int_{y_L}^{y_U} (d_A(y) - d_P(y))^2 dG(y)$$
$$+ \int_{y_U}^{\infty} (d_U - d_P(y))^2 dG(y) + s_x^2,$$

where y_L and y_U are defined by $d_A(y_L) = d_L$, $d_A(y_U) = d_U$. Since d_A is increasing we can again look for optimal values y_L^*, y_U^*. It will suffice to analyze the upper limit because the analysis of the lower limit is essentially identical.

Differentiating E_P w.r.t. y_U gives the first-order condition for y_U^*:

$$\int_{y_U^*}^{\infty} (a_A(y_U^* + m) + b_A - a_P(y + m) - b_P)dG(y) = 0.$$

With a Normal distribution this reduces to:

$$\text{(2.15)} \qquad \frac{a_A(y_U^* + m) + b_A - a_P \cdot m - b_P}{a_P} = \frac{g(y_U^*) \cdot s_y^2}{1 - G(y_U^*)} .$$

We notice that the function:

$$h(y) = \frac{g(y)}{1 - G(y)} ,$$

is the hazard rate for the Normal distribution; a well-known function from reliability theory (see Barlow and Proschan [1975]). In Appendix 2.B we show the following properties of the hazard rate:

$$\text{(2.16)} \qquad h(y) > y, \ \forall y;$$

$$\text{(2.17)} \qquad 0 < h'(y) \cdot s_y^2 < 1, \ \forall y;$$

(2.18) $$\lim_{y \to \infty} \frac{h(y)}{y} = \frac{1}{s_y^2} .$$

Thus the RHS in (2.15) has y as an asymptote. Graphically (2.15) can be described as follows:

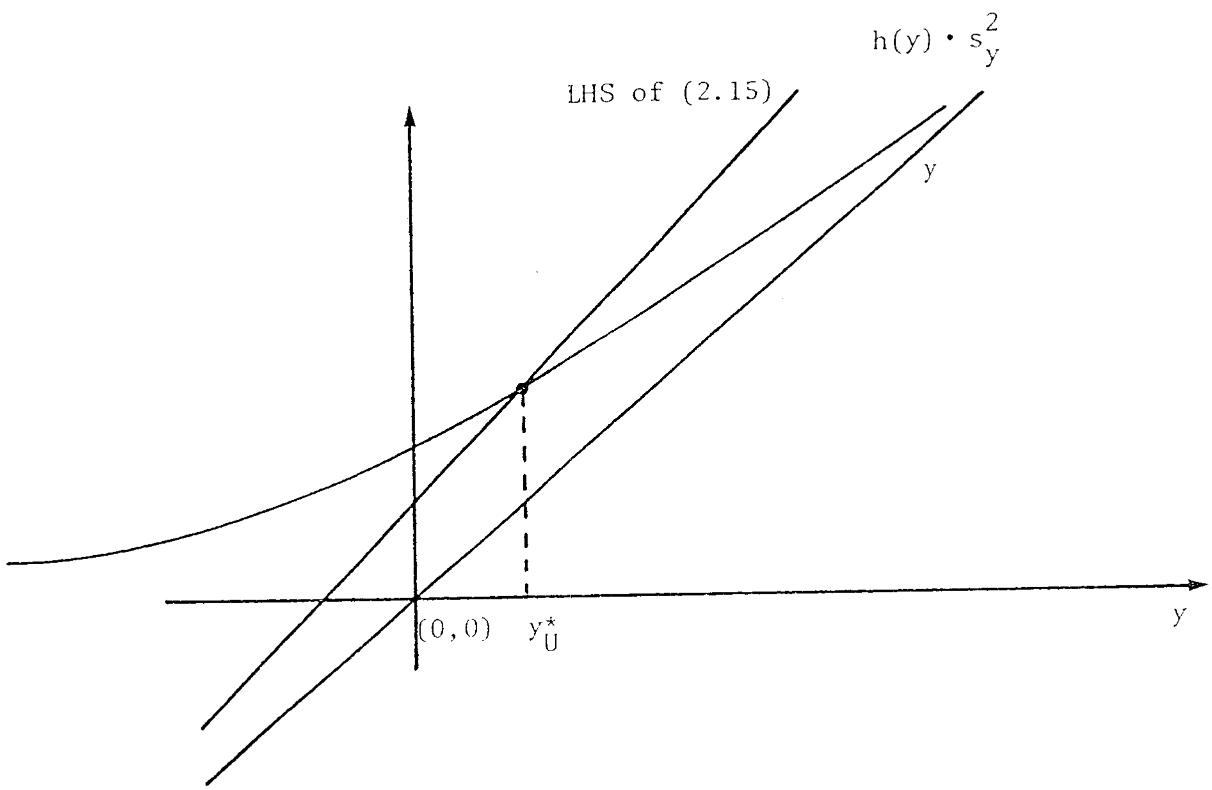

We assumed $a_A \geq a_P > 0$, $b_A \geq b_P$. Consequently, (2.16)-(2.18) imply that (2.15) has a unique solution $y_U^* < \infty$, unless $a_A = a_P$, $b_A = b_P$, in which case the principal's and the agent's response functions coincide and the agent is given complete freedom. Since $h(y) \cdot s_y^2 > 0$,

$$d_U^* = a_A(y_U^* + m) + b_A > a_P \cdot m + b_P = d^*,$$

where d* is the principal's best centralized act.[12] Similarly, we can show that $d_L^* < d^*$ and thus d* is contained in the optimal delegation interval (d_L^*, d_U^*).

It is readily checked that y_U^* corresponds to a global minimum of the expected loss. $a_A \geq a_P$ and (2.17) imply that the sign of the derivative changes from negative to positive at y_U^*, which guarantees a local minimum, and since y_U^* is a unique solution to (2.15) it is also global.

We summarize the preceding discussion in:

<u>Proposition 2.3</u>: With quadratic loss functions defined in (2.11) s.t. $a_A \geq a_P > 0$, $b_A \geq b_P > 0$, $(a_A, h_A) \neq (a_P, b_P)$, and with a Normally distributed information gap $\tilde{y}$, the optimal interval control is nondegenerate and finite and contains the best centralized act.

Now look at changes in the information gap $\tilde{y}$. Only the RHS of (2.15) changes with s_y^2. In Appendix 2.B we show that,

$$(2.19) \qquad \frac{\partial}{\partial s_y}\,[h(y) \cdot s_y^2] > 0 \qquad , \text{ for all } y.$$

Since a_A, $a_P > 0$, it follows from (2.15) (by total differentiation w.r.t. s_y) that $\frac{\partial y_U^*}{\partial s_y} > 0$, which implies $\frac{\partial d_U^*}{\partial s_y} > 0$. A symmetric analysis shows that $\frac{\partial d_L^*}{\partial s_y} < 0$. We have thus:

Proposition 2.4: Under the assumptions of proposition 2.3, the agent is given more freedom with an increase in the information gap.

As we said before, the change in the information gap can be interpreted either so that the principal faces more uncertainty or the agent gets more informed.

The implications of changes in preference parameters will be given an exhaustive treatment in Section 2.3.6, so we will not make any general conclusions from (2.15). For the particular loss functions in Example 2.2 it is, however, immediate from (2.15) and (2.17) that a decrease in $(b_A - b_P)$ will increase y_U^* and hence d_U^*. In other words, when preferences are closer, the agent has more freedom. In conjunction with Proposition 2.4, this shows that the same conclusions we made about changes in information and preferences when $\tilde{y}$ was uniformly distributed (see (2.10)), are true when $\tilde{y}$ has a Normal distribution.

Suppose the agent gets more informed while the principal's information stays the same. Does the principal's welfare increase with this change? The answer is not generally in the affirmative as we will see later, but for the preference functions in Example 2.2 the statement is true. The argument is simple and we will omit detailed calculations.

Let the agent be able to observe either y_1 or $y_1 + y_2$. The principal should ask himself the question: conditional on y_1, would he like the agent to observe also y_2 and make a corresponding change in the decision? When response functions are parallel as in Example 2.2, the agent has already taken out his bias after he makes a decision based on y_1. The change in his decision that y_2 will induce will only reduce the risk for the principal. Moreover, the upper limit will add to the benefit, since if $d_A(y_1) > d_U$ so that the agent is constrained to d_U, then y_2 can only cause him to decrease his response, which always is beneficial for the principal.

A graphical illustration will make the point clear:

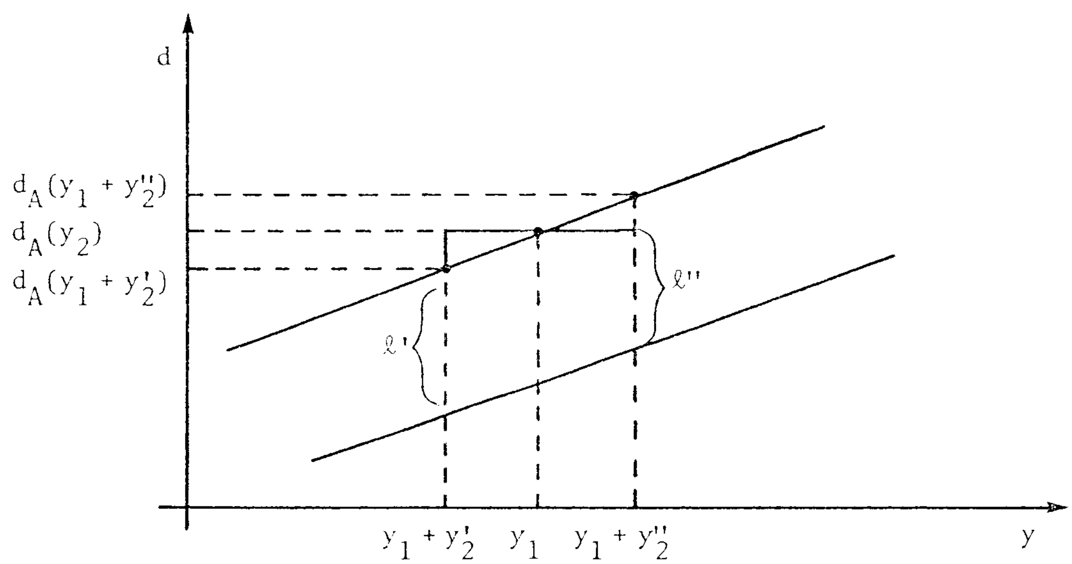

y_1 is observed and y_2', y_2'' are two tentative outcomes for y_2 that lie symmetrically around y_1. Because the principal's loss function is quadratic with the same curvature for all y and $\ell'' < \ell'$ in the figure above, the loss from taking $d_A(y_1 + y_2'')$ rather than $d_A(y_1)$ at $y_1 + y_2''$, is less than the loss from taking $d_A(y_2)$ rather than $d_A(y_1 + y_2')$ at $y_1 + y_2'$. This difference will only be enhanced if $d_A(y_1 + y_2'') > d_U$ so that the agent gets constrained at d_U. Consequently, letting the agent observe $y_1 + y_2$ rather than y_1 alone will be preferred by the principal.

This argument can be made algebraically and applies also for the more general loss functions in (2.11) as long as $a_A < 2a_P$; (this is the point where ℓ'' becomes greater than ℓ' in the picture above). Though we conjecture that the statement is true for all loss functions of the form (2.11), we do not have a general proof. The complication arises because with $a_A > 2a_P$ the role of the upper limit becomes essential and we do not have an explicit expression for it; (with complete freedom, i.e., $d_L = 0$, $d_U = +\infty$, the principal is better off with a more informed agent only if $a_A < 2a_P$).

Of course, the quadratic loss functions we have used in this example (see (2.11)) are quite special. But they give an indication of how quantity controls can be employed for improved decision-making and what determines the optimal amount of delegation. We can also argue that quadratic loss functions represent a first-order approximation of more general loss functions.

Before embarking on an analysis of how far the results of the example generalize, we will apply the same quadratic model to a study of quantity controls in economic planning.

2.3.3 Quantity Controls in Economic Planning

The literature on economic planning has focused predominantly on iterative planning models for reaching efficient allocations (see Heal [1973]). This literature owes much to the theory of mathematical programming, from which various algorithms have been adopted and applied to planning problems. One of the rewards of this line of research has been the appealing interpretations of decentralized planning procedures that can be given to most of the algorithms (see Jennergren [1971]). To some extent one can also use the convergence properties of these algorithms as guidelines for designs of practical planning procedures.

However, many interesting issues cannot be successfully attacked by this approach. One of the shortcomings is that little can be said about what should be done if the iterative exchange of information has to be stopped short of reaching an optimal solution. And yet this is clearly the common situation in practice. Maybe only one or two iterations will normally be carried out before a final decision has to be made, and in that case hardly anything can be learned from studying planning algorithms.

In an insightful article, Weitzman [1974] approaches this question from an entirely new perspective. He recognizes that when an information gap between economic agents and the center remains at the time of decision-making, it matters which instrument of implementation is used. More specifically, he asks: What is the best mode of implementation -- prices or quantities? This should be contrasted to the classical result that it is a matter of indifference which of the two modes is used, once the center has received the necessary information for optimal decision-making.[13]

We will start by describing briefly Weitzman's analysis of the question, including an important comment by Laffont [1977]. Our main interest lies in extending the analysis to mixed schemes where agents face both price and quantity controls. Such schemes are superior to either prices or quantities alone. Moreover, studying them will aid our understanding of how the tightness of economic control depends on both the information gap between the center and the agents, and the structure of preference functions. This stands in contrast to the results of Weitzman's analysis, which indicates that the superiority of price or quantity controls is solely determined by the preference functions.

Prices vs. Quantities

Weitzman analyzes a simple model with one commodity of which an amount d can be produced at a cost $C(d,z)$ and results in benefits $B(d)$; z is the state of nature. It is assumed that $C''(d,z)$, $B''(d) < 0$, $B'(0) > C'(0,z)$ for every z, and for every z, $B'(d) < C'(d,z)$ for d sufficiently large. The production unit knows z; the center knows $B(d)$. The center's problem is to decide on two alternative modes of operation:

I. Announce a quantity d^* to be produced with no other choice for the production unit.

II. Announce a price p^* and let the production unit determine d by maximizing profits $p^* \cdot d - C(d,z)$.

When Option I is used, the optimal quantity d^* is determined from the first-order condition. Weitzman asks under what conditions one or the other control mode is preferred.

$$(2.20) \qquad B'(d) = E(C_1(d,z)),$$

where C_1 is the partial derivative w.r.t. d. In this case the production unit's superior information about costs is entirely ignored.

When a price p is announced (Option II), the producer maximizes his profits, which results in a response function $d(z;p)$ that satisfies:

$$(2.21) \qquad p = C_1(d(z,p),z).$$

The optimal price announcement p^* maximizes:

$$B(d(z,p)) - E[C(d(z,p),z)].$$

From the first-order condition we get, using (2.20):

$$(2.22) \qquad p^* = \frac{E[B_1(d(z,p^*)) \cdot d_2(z,p^*)]}{E[d_2(z,p^*)]} ,$$

where d_2 is the partial derivative w.r.t. p. This condition determines p^*.

The comparative advantage of prices over quantities is determined naturally by:

$$(2.23) \qquad \Delta \equiv E[B(d(z,p^*)) - C(d(z,p^*),z)] - E[B(d^*) - C(d^*,z)].$$

If $\Delta > 0$, prices are preferred, otherwise quantities. In order to arrive at an explicit expression for Δ, Weitzman makes the following quadratic approximations of B and C in the neighborhood of d^*.[14]

$$(2.24) \qquad B(d) = b + B' \cdot (d-d^*) + \frac{1}{2} \cdot B'' \cdot (d-d^*)^2,$$

$$(2.25) \qquad C(d,z) = c(z) + (C' - h(z)) \cdot (d-d^*) + \frac{1}{2} \cdot C'' \cdot (d-d^*)^2.$$

W.ℓ.o.g. we can assume $h(z) = z$, and $E_z = 0$, because $c(z)$ will play no role in the analysis. In that case, z measures marginal costs with an increase in z corresponding to a decrease in marginal costs. In (2.24)-(2.25), B', C', B'', C'' are constants with $B'' < 0$, $C'' > 0$.

From (2.20) it follows that $B' = C'$. The producer's response function when a price p is used will be:

(2.26) $$d(z,p) = d^* + \frac{p - C' + z}{C''} .$$

$d_2(z,p^*) = \frac{1}{C''}$, and substituting this into (2.22) gives:

(2.27) $$p^* = B' = C'.$$

Some simple substitutions finally produce the result:

(2.28) $$\Delta = \frac{s^2 \cdot (B'' + C'')}{2 \cdot C''^2} ,$$

where $s^2 = Ez^2$. This formula is the main result in Weitzman's paper. We can see that the sign of $B'' + C''$ alone determines which control mode should be used. When $-B'' < C''$ prices are preferred; otherwise quantities. This is intuitive. $-B'' >> 0$ means that the benefit function is relatively curved at d^*, and so small changes in d will result in large changes in benefits. In this situation the center does not want to take the risk of announcing a price and have d vary with the producer's cost function (the expected variation being larger with small C''), and consequently, prefers direct control of d. On the other hand, when $B'' = 0$ the benefit function is linear and so a price will communicate all the center's information to the producer and a first-best optimum can be achieved.

Less intuitive is the result that the information gap between the center and the production unit, measured by s^2, does not play any role in determining the optimal control mode. The information gap only magnifies the comparative advantage of the superior mode. From this we may make the inference that the tightness of economic control, at least as far as the two extreme control modes are concerned, does not depend on how much information the center possesses. That this statement is false when intermediate control modes can be used will be shown shortly. We will verify the more intuitive conjecture that the less information the center has, the more it will tend toward price-guided control.

Insightful as Weitzman's article is, it has been a source of inspiration for many subsequent papers.[15] Laffont has observed that Weitzman's analysis is one-sided in that only the control of suppliers is considered. The center represents the demanders' side. For some economic activities this is the appropriate model; for instance, when analyzing pollution control or corresponding public goods problems. On the other hand, with private goods the center could possibly use prices on the demand side rather than the supply side as a control instrument, in order to utilize the demanders' superior information about benefits.[16]

For this reason, assume that the benefit function is $B(d,z_1)$, where z_1 is only known to the demander. In the preceding analysis we had no reason to include this uncertainty, since only the expectation of the benefits was of relevance. As before, the cost function is $C(d,z_2)$, with z_2 only known to the supplier. Three options are now available for the center:

I. Order d* from the supplier and announce d* as the consumption to the demander.

II. Announce a price p* to the supplier. Let him choose that d which maximizes his profits and announce this as the consumption to the demander.

III. Announce a price p* to the demander, and order the amount d he demands from the supplier.

The third mode is new. Using the same approximations as before (with obvious modifications for B), Laffont derives the following comparative advantages between the three modes:

$$\Delta(II/I) = \frac{s_2^2(B'' + C'')}{2C''^2}, \tag{2.29}$$

$$\Delta(III/I) = \frac{-s_1^2(B'' + C'')}{2B''^2}, \tag{2.30}$$

$$\Delta(II/III) = \left[\frac{s_1^2}{2C''^2} + \frac{s_2^2}{2B''^2}\right] (B'' + C''). \tag{2.31}$$

Here $\Delta(II/I)$ is the comparative advantage of Option II over Option I, etc.; $s_1^2 = Ez_1^2$, $s_2^2 = Ez_2^2$.

The important observation is that everything in favor of quantity orders in Weitzman's analysis is generally in favor of using prices on the consumption side. Indeed, I is never optimal, so the choice is between using prices for either the producer or the consumer. As Laffont points out, this result depends on the quadratic approximations used. With the second-order term random in (2.24)-(2.25), quantities may be preferred over both price schemes. Regardless of this, Laffont's analysis points out the fundamental duality between price controls in the two sectors of the economy.

In the context of a firm, Laffont's model applies to the transfer pricing problem. From (2.31) we can get a tentative answer to the question: which department should be given the right to determine the amount of transfer? The answer is: the department with a more curved benefit or cost function, that is, the department for which variation in the amount transferred carries a higher cost. This can also be seen easily from a Marshallian supply-demand diagram.

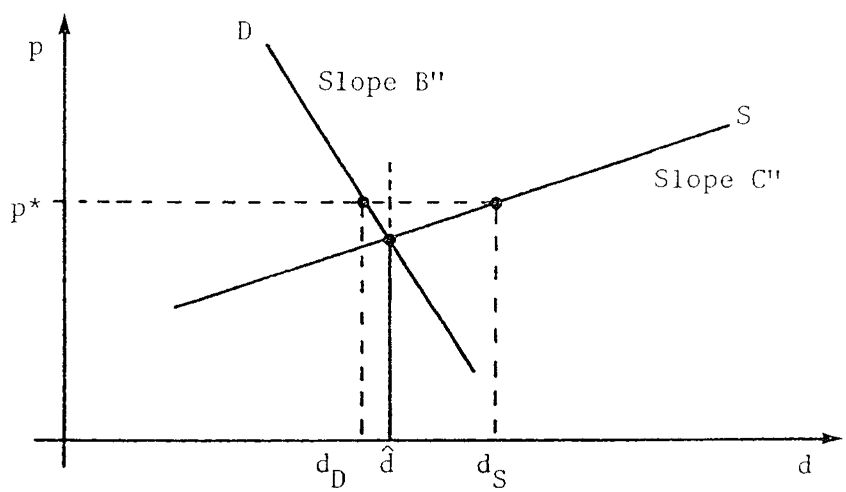

From the quadratic form of the benefit and cost functions we get linear demand and supply curves. The slopes are B''(<0) and C'' respectively. The random variables z_1 and z_2 only cause a vertical shift in these curves. In the picture above we have drawn a particular outcome of (z_1, z_2), and we can see that if $-B'' > C''$, less losses (in terms of consumer and producer surpluses) are incurred when d_D is taken than when d_S is taken, regardless of what price is used. Hence, the demander's wishes should be followed in this situation. The opposite is true if $C'' > -B''$.

In passing we can compare Laffont's solution to a standard market outcome in disequilibrium.[17] When the price is fixed, but supply and demand differ, the long side will be constrained, i.e., the outcome will be $\min(d_D, d_S)$. Such a control mechanism is also quite natural in a centrally planned economy. The interesting point is that, with linear demand and supply functions, announcing the same price to both sides of the market and transferring the minimum of demand and supply, will always produce an inferior solution to either Option II or Option III in Laffont's scheme. Thus, a disequilibrium solution can be dominated by letting one side of the market determine the outcome alone.

Quantity Controls

A natural extension of Weitzman's analysis is to look at a mixed scheme where a price p^* is announced, and in addition to this, upper and lower limits on the quantity are imposed. Let (d_L, d_U) be the range in which the producer has freedom to choose d, when the price p^* is given. As extreme cases we have, of course, $d_L = d_U = d^*$ and $d_L = 0$, $d_U = \infty$, which correspond to the use of a quantity order and to a pure price scheme. First, we will show that a quantity order can always be dominated by using a mixed scheme $[p, (d_L, d_U)]$. In other words, some freedom should always be given to the production unit.

<u>Proposition 2.3</u>: If d* is the best centralized act,[12] then there exists a nondegenerate interval (d_L, d_U) containing d*, and a price p, such that both the center and the production unit are better off with this control than with d*.

<u>Proof</u>: Let $d_A(z,p)$ be the production unit's response function when a price p is announced, and $d_p(z)$ the center's response function; i.e.,

(2.32) $$d_A(z,p) = \underset{d}{\operatorname{argmax}} [p \cdot d - C(d,z)],$$

(2.33) $$d_p(z) = \underset{d}{\operatorname{argmax}} [B(d) - C(d,z)].$$

Let d(p) be the solution to $B'(d) = p$. By our convexity assumptions,

(2.34) $$\begin{cases} d_A(z,p) \leq d(p) \Rightarrow d_p(z) \geq d_A(z,p), \\ d_A(z,p) > d(p) \Rightarrow d_p(z) < d_A(z,p). \end{cases}$$

In other words, the production unit makes a too-small production decision whenever its decision is below d(p), and vice versa when it is above d(p) (see picture below).

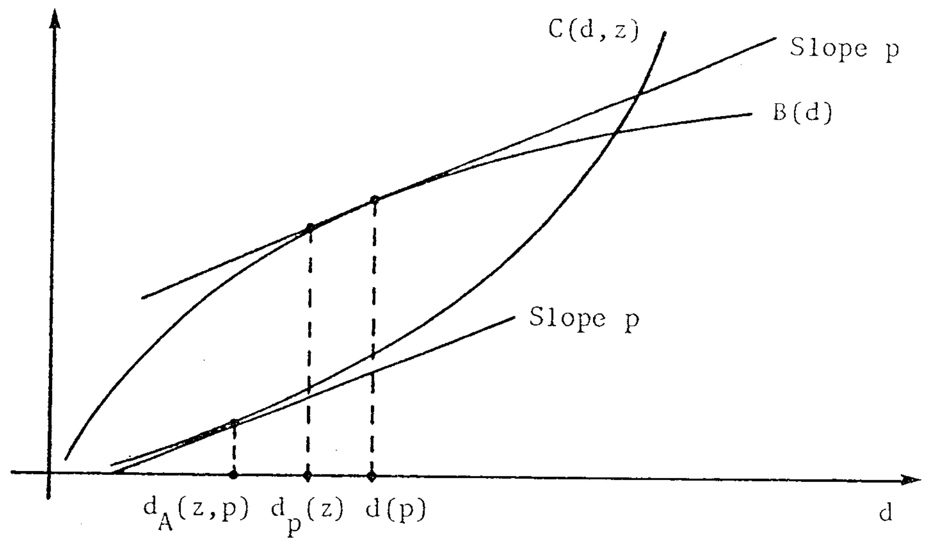

Now, set the price $\hat{p}$ such that $d(\hat{p}) < d^*$, and take $d_L = d(\hat{p})$, $d_U = d^*$. Whenever the production unit decides on a d inside the control interval (d_L, d_U), the center will strictly prefer this decision to d^*, by (2.34) and our convexity assumptions. By definition, the production unit will prefer its own choice to d^*. Since this is true for every z, both parties will be better off with the interval control $(d(\hat{p}), d^*)$ and the price $\hat{p}$.

Q.E.D.

It is of interest to ask when a pure price scheme can be dominated in a similar way by a price-quantity control pair. Of course, a necessary condition is that B is not linear, which we already assumed. A sufficient condition is that for every price p, there is an upper limit for the amount the production unit will produce; i.e., sup $d_A(z,p) < \infty$. This will become clear from the forthcoming discussion and we will not pursue the topic further here. Instead, we will turn to quadratic approximations of the benefit and cost functions in order to analyze properties of the optimal quantity control.

Assume (2.24)-(2.25) are valid approximations of B and C, when z is Normally distributed with mean = 0 and variance = s^2. If d* is large enough compared to s^2, we can ignore the fact that the producer, according to his response function (2.26), will make negative production decisions. We could use a truncated Normal distribution, but this would lead to unnecessarily complicated calculations.

The center's response function, if it knew z, is:

$$d_P(z) = d^* + \frac{z}{C'' - B''}, \qquad (2.35)$$

and the producer's is

$$d_A(z,p) = d^* + \frac{p - C' + z}{C''}. \qquad (2.36)$$

These are depicted in the figure below.

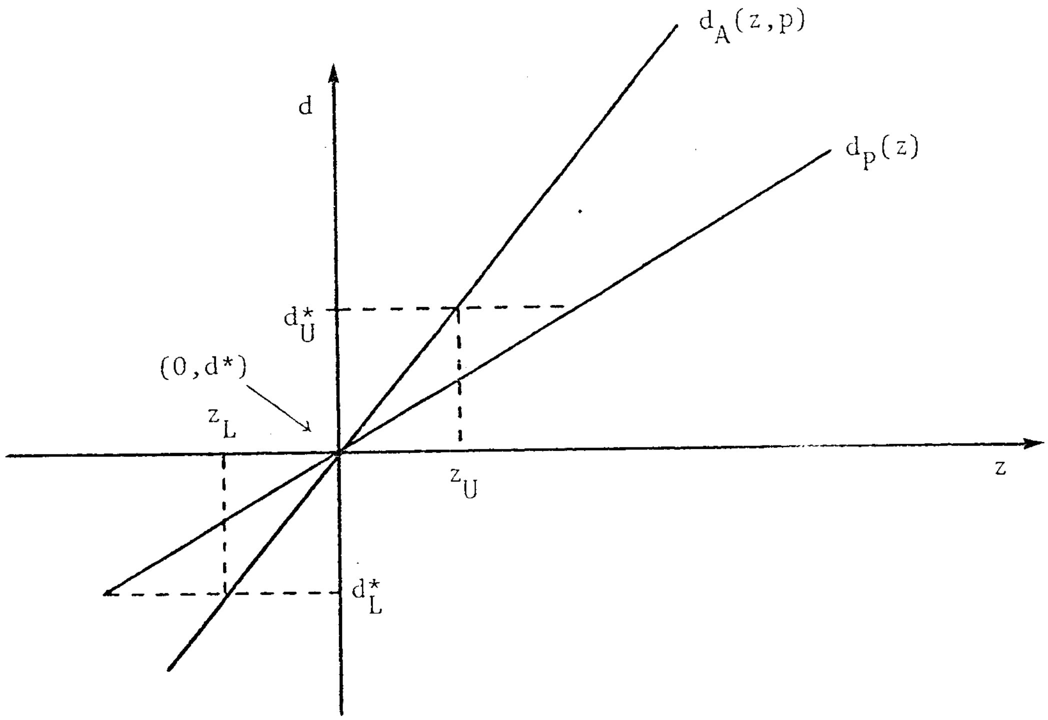

Since C" - B" > C", the production unit's response function is steeper than the center's. The slope is independent of the price, which only affects the constant term. If the optimal price $p^* = B' = C'$ is used, it is clear by symmetry of the response functions and the distribution of z, that for the optimal quantity control (d_L^*, d_U^*), $d_U^* - d^* = d^* - d_L$. Given these control limits it also follows by symmetry of the distribution and the quadratic approximation that p* is the optimal price to use.

We find that the center's objective function B(d) - C(d,z) and the production unit's objective function $p^* \cdot d - C(d,z)$, when approximated as in (2.24)-(2.25), are special cases of the quadratic loss functions in (2.11). The identification is $b_A = b_P = d^*$, $a_A = 1/C''$, $a_P = 1/(C'' - B'')$. Since $B'' < 0$, we have $a_A > a_P$ as was assumed in our analysis. We can then apply directly the results of the previous section here. We conclude:

Proposition 2.5: With a Normally distributed information gap $\tilde{z}$:

(i) the production unit will be given a finite degree of freedom which will include the best centralized act d^*. Neither pure prices nor direct quantity orders are optimal;

(ii) the production unit is given more freedom with a decrease in the curvature of the benefit function or an increase in the curvature of the cost function;

(iii) the production unit will be given more freedom with an increase in the information gap $\tilde{z}$.

Proof: (i) is equivalent to Proposition 2.3. (ii) follows from (2.15) and (2.17), since $m = 0$ and $b_A - b_P = 0$. (iii) is equivalent to Proposition 2.4. Q.E.D.

Remark: Even though we have assumed that the production unit knows z (in congruence with Weitzman's model), the production unit could have received an imperfect signal, $\tilde{y}$. This would allow us to interpret changes in information as before, both in terms of the center being less informed or the production unit acquiring more information.

Part (ii) of Proposition 2.5 confirms Weitzman's result. When $|B''|$ increases or C'' decreases, the response functions in (2.35) and (2.36) diverge, increasing the need for more rigid control. Part (iii), however, is distinctly different from Weitzman's formula. We can say slightly more about the behavior of d_U^* as the information gap increases. We have:

$$(2.37) \qquad \frac{g(0) \cdot s^2}{1 - G(0)} = 2 \cdot \psi(0) \cdot s \to \infty, \text{ as } s \to \infty,$$

where g and G are the density and distribution functions of z, and ψ is the density function of the standardized Normal distribution. Since the hazard rate is increasing by (2.17), (2.37) implies that the RHS of (2.15) goes pointwise to infinity with s for positive values of the argument. Hence, $d_U^* \to \infty$, as $s \to \infty$, which means that the bigger the information gap gets, the closer we get to the pure price mechanism. This stands in contrast to Weitzman's formula (2.28), which says that if $B'' + C'' < 0$, so that quantities are preferred, then the expected loss from using prices rather than quantity orders goes to infinity with increased information gap.[18] No contradiction is involved, however. Our result merely points out that one should be careful with interpreting Weitzman's formula in terms of tightness of economic control. Even though quantity orders may be "infinitely better" than prices because of large variance in $\tilde{z}$, the optimal quantity control may give the economic agents a large interval of freedom.

We can take (iii) above as an indication that the tightness of economic control also depends on how informed the center is about the economic conditions of the agents. The less informed the center is, the less it should intrude in the economic activities via rigid quantity constraints.

The analysis has only been carried out for the supply side of the market, but the results are analogous on the demand side. Hence, Laffont's scheme can likewise be improved by using quantity controls in addition to prices.

We turn now to a study of general preference functions and distributions in a principal-agent framework.

2.3.4 Interval Controls

We have worked exclusively with interval controls in the preceding analysis, and it is natural to ask under what conditions the overall optimal quantity control is indeed an interval. We saw that this was the case in the production problem (Example 2.2), when $\tilde{y}$ had a uniform distribution.

Despite the wide use of interval controls, it does not appear easy to give natural conditions under which they are optimal in the model we have presented. Before indicating the problems, it should be pointed out that there may be other reasons why intervals are used. Clearly, they are simple to use with a minimal amount of information and monitoring needed to enforce them. Under certain circumstances the only relevant controls are intervals because of the enforcement problem. A case in point is bank lending. Even if banks used some type of discontinuous controls with certain loan ranges prohibited, the borrower could get around this constraint either by using several banks or by borrowing in excess of his needs and reinvesting the extra money in riskless assets.

Let us study only increasing response functions and concave preferences. In the picture below we have drawn a control with a gap.

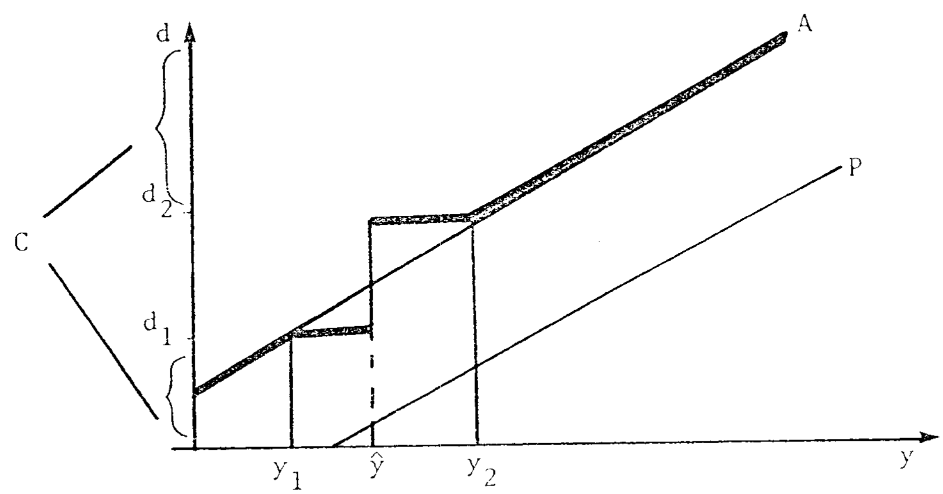

The agent is not allowed to pick a $d \varepsilon (d_1, d_2)$. If $y \varepsilon (y_1, \hat{y})$ he will choose d_1, and if $y \varepsilon (\hat{y}, y_2)$ he will choose d_2. From the principal's point of view, the agent's modified response means an improvement in the region $(y_1, \hat{y})$ but a worse decision in $(\hat{y}, y_2)$. The net benefit depends on the distribution of $\tilde{y}$, on where $\hat{y}$ is located and on the particular form of the principal's preference function. This should indicate the problems one faces in trying to give general conditions under which the interval rule is optimal.

Very special and simple extensions of the production problem (Example 2.2) can be given. For instance, if the loss functions are symmetric around the optimal response and independent of y (or become uniformly steeper as y increases), and if the response functions are linear and diverging, then the interval rule is overall optimal for a uniformly distributed $\tilde{y}$. It seems intuitively plausible though, that if the principal is sufficiently uncertain about the response function of the agent, he will not know where to leave a gap in the control, and this would give rise to optimality of the interval rule.

In the next sections on one-dimensional controls, we will continue to restrict attention to interval controls, regardless of their overall optimality. If the optimal control has gaps, it is quite hard to derive, and certainly an analysis of interval controls can be defended on pragmatic grounds.

In our examples we have found that the best centralized act d* is always included in the optimal control interval. This is generally true, provided a coherence condition is met.

Define the <u>level sets</u>,

$$(2.38) \qquad Y_i(d) = \{y \varepsilon Y | d_i(y) \geq d\}, \qquad i = A, P,$$

of an agent A and the principal P. $Y_A(d)$ is the set of signal outcomes under which the agent would prefer a higher action than d. Obviously, $Y_i(d) \subseteq Y_i(d')$, when $d' \leq d$.

<u>Definition</u>: The agent's and the principal's preferences (or preference functions) are said to be <u>coherent</u>, if for every $d \in D$:

$$Y_A(d) \cap Y_P(d) = Y_A(d) \text{ or } Y_P(d).$$

To understand this definition better, let us look at a violation of coherence.

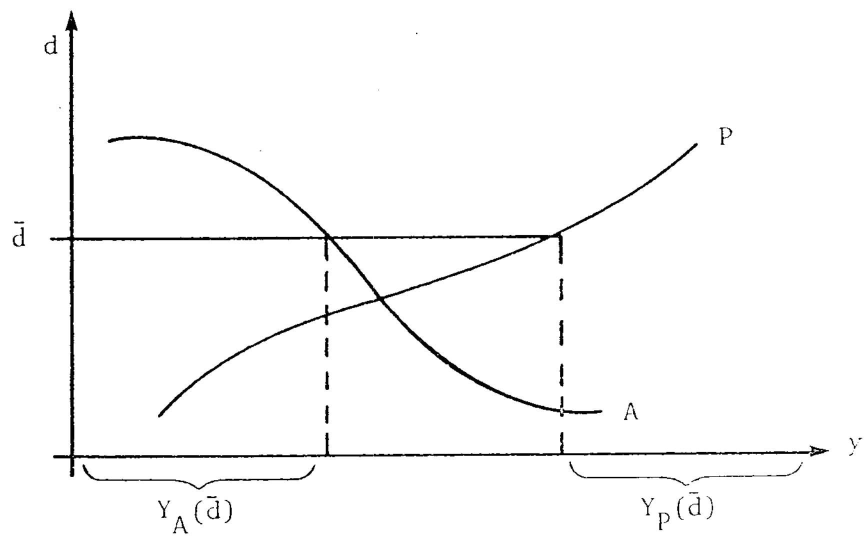

In the situation pictured above the agent's behavior is in direct conflict with the principal's in terms of how he reacts to y. In fact, the principal can have no use of the agent as we will argue later. Even if this case is extreme, the value of decentralization via quantity controls is in doubt when preferences are not coherent. This is also evident from rephrasing the definition: if preferences are coherent, there cannot be two signals such that the agent takes the same action for both signals, but the principal prefers a higher action than the agent's for one of the signals and a lower action for the other. (If there were two such signals y_1 and y_2, then for some d, $y_1 \varepsilon Y_P(d)$, $y_1 \not\varepsilon Y_A(d)$ and $y_2 \varepsilon Y_A(d)$, $y_2 \varepsilon Y_P(d)$, contradicting coherence).

In all of the examples we have discussed the preferences have been coherent, since coherence is implied by increasing response functions. It also holds whenever the agent is constantly biased toward either too high or too low actions. Notice that coherence is not related to the concavity of preferences; it is a property of response functions alone. In the sequel, we will only deal with coherent preferences.

Theorem 2.6: Assume that preferences are concave, coherent and continuous. Then there exists an optimal interval control, which contains the principal's best centralized decision.[19]

Proof: Let the optimal interval control [a,b] lie above the principal's best decision d*. The other possibility can be treated similarily. Let us furthermore assume [a,b] is an optimal interval with the lowest lower bound. Such an interval exists by our continuity assumption, and because the set of optimal lower bounds above d* is closed. We will show that there exists an $\varepsilon > 0$ s.t. $[a - \varepsilon, b]$ is at least as good as [a,b], unless a = d*. This will prove our claim.

Suppose $a > d^*$, contrary to the assertion in the theorem. By coherence, two possibilities exist: (i) $Y_P(a) \subseteq Y_A(a)$, (ii) $Y_A(a) \subseteq Y_P(a)$. They are pictured in the figures below.

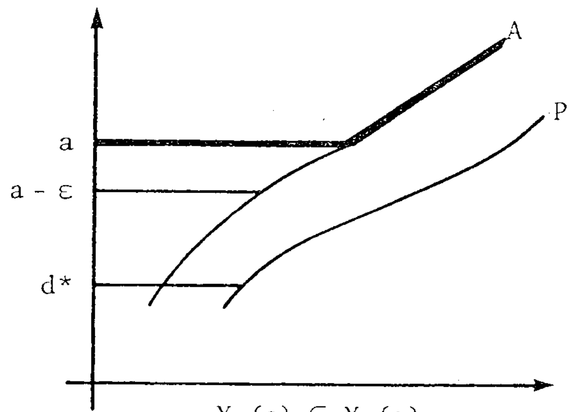

$Y_P(a) \subseteq Y_A(a)$

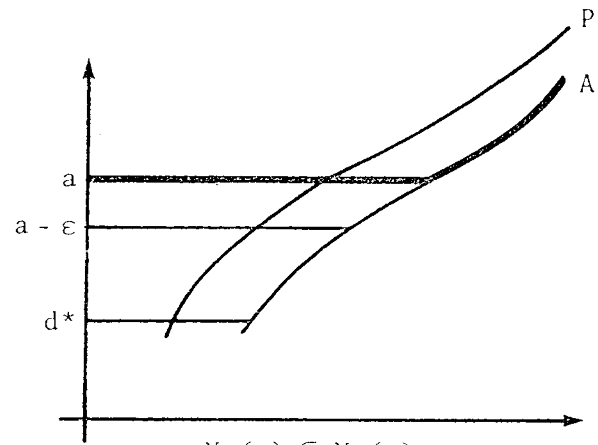

$Y_A(a) \subseteq Y_P(a)$

It could be that $Y_P(a) = Y_A(a)$. We will for the moment assume this is not the case, and come back to it later.

<u>Case 1</u>: $Y_P(a) \subseteq Y_A(a)$, $Y_P(a) \neq Y_A(a)$. By coherence and continuity of preferences, there exists an $\varepsilon > 0$ s.t. $d_P(y) \leq d_A(y)$, for $y \varepsilon Y_A^C(a - \varepsilon) - Y_A(a)$, and $d_P(y) \leq a - \varepsilon$, for $y \varepsilon Y_A^C(a - \varepsilon)$. By concavity of the agent's preference function,

$$a - \varepsilon \leq d_A(y|[a - \varepsilon, b]) \leq d_A(y|[a,b]), \text{ for } y \varepsilon Y_A^C(a),$$

and

$$d_A(y|[a - \varepsilon, b]) = d_A(y|[a,b]), \text{ for } y \varepsilon Y_A(a).$$

It follows that under the control $[a - \varepsilon, b]$ the agent's response will be pointwise as close to the principal's as under the control $[a,b]$. Since the principal's preference function is concave, $[a - \varepsilon, b]$ is at least as good as $[a,b]$, contradicting the minimality of $\underline{a}$.

Case 2: $Y_A(a) \subseteq Y_P(a)$, $Y_A(a) \neq Y_P(a)$. In this case, when we lower the bound to $a - \varepsilon$, it is not true that the agent's response will become pointwise closer (see figure). We have to argue differently. By continuity and coherence of preferences there exists an $\varepsilon > 0$ s.t.

$$d_P(y) \geq d_A(y), \qquad \text{for } y \varepsilon Y_A(a - \varepsilon) - Y_A(a). \tag{2.39}$$

Furthermore $Y_A(a) \subseteq Y_P(a)$ implies

$$d_P(y) \geq a, \qquad \text{for } y \varepsilon Y_A(a). \tag{2.40}$$

By concavity of the principal's preference function, he considers $d = a - \varepsilon$ at least as good a constant response as $d = a$ (since the integral of a pointwise concave function is concave and $a - \varepsilon$ is closer to d^* than $\underline{a}$). It follows that the principal prefers $d = a - \varepsilon$ to $d = a$ on $Y_A^c(a)$, since he prefers $d = a$ to $d = a - \varepsilon$ on $Y_A(a)$ by (2.40). By (2.39) the principal prefers (weakly) $d_A(y|[a - \varepsilon, b])$ to $d = a - \varepsilon$ on $Y_A^c(a)$, since the former is pointwise closer in this region. Combined we get that the principal finds $d_A(y|[a - \varepsilon, b])$ at least as good as $d_A(y|[a,b])$ on $Y_A^c(a)$. On $Y_A(a)$

these responses coincide so we have shown that $[a-\varepsilon, b]$ is at least as good as $[a,b]$ contradicting the minimality of $\underline{a}$.

Case 3: $Y_A(a) = Y_P(a)$. By coherence of preferences we get either into Case 1 or Case 2 above, when we lower the bound to $a-\varepsilon$, and we get a contradiction as before.

Consequently, $a > d^*$ is not possible. An analogous proof shows $b < d^*$ is not possible and the theorem is proved. Q.E.D.

We want to stress that nothing was assumed about the dimensionality of $\tilde{y}$ in the theorem. With coherent preferences most results that are true for a one-dimensional $\tilde{y}$ and increasing response functions, will carry over to the general case. This will be seen repeatedly in the sequel.

Notice that, since the agent is better off when he is given an extended set of choices, Theorem 2.6 implies that whenever it is optimal for the principal to use a nondegenerate interval control, it is also beneficial for the agent. In other words, a Pareto improvement (from the best centralized act) is guaranteed if it pays to decentralize.

2.3.5 Changes in Information

The fact that the agent has some private information is the driving force behind decentralization. From this statement it is natural to conjecture that the "more" private information the agent has, the more he should be given freedom. This was seen to hold true for the examples in Sections 2.3.2 and 2.3.3. In general it is not true. As a counterexample we can look at the following discrete case:

Example 2.6: $\tilde{z}$ is uniform on (0,1). In the coarser information system the agent has a partition on $\tilde{z}$ which is {(0,1/2), (1/2,3/4), (3/4,1)}; in the finer information system the partition is {(0,1/4), (1/4,1/2), (1/2,3/4),(3/4,1)}. The principal's partition is {(0,1)}. $D = \{d_1,d_2,d_3\}$. The preferences are given in the table below:

	$z\varepsilon(0,1/4)$	$z\varepsilon(1/4,1/2)$	$z\varepsilon(1/2,3/4)$	$z\varepsilon(3/4,1)$
d_1	(10,10)	(10,5)	(0,0)	(0,0)
d_2	(0,0)	(0,0)	(10,10)	(9,9)
d_3	(0,0)	(0,10)	(0,0)	(10,10)

The first number in each parenthesis indicates the principal's utility index, the second the agent's. When the agent has the coarser information system, he will choose d_1 if $z\varepsilon(0,1/2)$, d_2 if $z\varepsilon(1/2,3/4)$, and d_3 if $z\varepsilon(3/4,1)$, which is exactly how the principal would decide too. Hence, the optimal control is $C = D$. When the agent has the finer information system, he will change his previous response if $z\varepsilon(1/4,1/2)$ from d_1 to d_3. This lowers the principal's expected utility by 2.5 if $C = D$, and since the advantage of having d_3 in the control set was small before, the optimal control with the fine information system is $C = \{d_1,d_2\}$.

The conclusion is that the principal gets worse off and gives the agent less freedom as the agent gets more informed. []

For continuous response functions similar examples could be given. For instance, the agent's loss function may be asymmetric so that when he faces some uncertainty he will decide as the principal, but with the uncertainty removed he will not. This effect can be made strong enough to offset the benefits from reduced uncertainty.

A general discussion of the case where the agent gets more informed is difficult because of the aforementioned feature that his response function may also change. It is easier to look at the case where the principal gets less informed (or more generally, changes his beliefs about y), but the agent maintains his information structure. As we mentioned before, the two cases are equivalent for the quadratic loss functions defined in (2.11).

We will only look at a one-dimensional signal $\tilde{y}$ with increasing response functions. An illustration is given below.

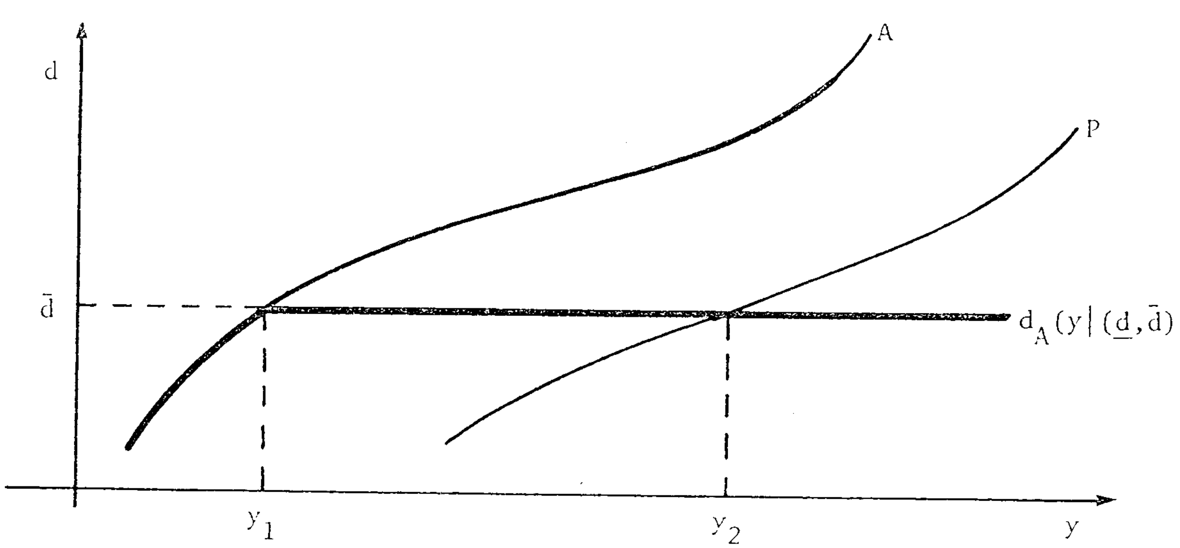

Let $f(y;\varepsilon)$ be a parameterized family of density functions, and $(\underline{d},\bar{d})$ the optimal control when $\varepsilon = 0$. We are interested in changes in the optimal control when ε increases.

Theorem 2.7: The agent is given more freedom with a differential change in the distribution of $\tilde{y}$, if and only if

(2.41)
$$-\int_{y_1}^{y_2} \frac{\partial}{\partial d} F^P(\bar{d},y) \frac{\partial}{\partial \varepsilon} f(y;0)\,dy < \int_{y_2}^{+\infty} \frac{\partial}{\partial d} F^P(\bar{d},y) \frac{\partial}{\partial \varepsilon} f(y;0)\,dy$$

and

(2.42)
$$-\int_{y_3}^{y_4} \frac{\partial}{\partial d} F^P(\underline{d},y) \frac{\partial}{\partial \varepsilon} f(y;0)\,dy > \int_{-\infty}^{y_3} \frac{\partial}{\partial d} F^P(\underline{d},y) \frac{\partial}{\partial \varepsilon} f(y:0)\,dy$$

assuming these differentials exist. Here

$$y_1 = d_A^{-1}(\bar{d}),\ y_2 = d_P^{-1}(\bar{d}),\ y_3 = d_P^{-1}(\underline{d}),\ y_4 = d_A^{-1}(\underline{d})$$

Proof: We only prove (2.41). Since the response functions are increasing we have a situation as pictured above, where $\bar{d}$ intersects d_A and d_P only once. Hence, the principal decides on the optimal interval limit $\bar{d}(\varepsilon)$ by solving d from

$$\frac{\partial}{\partial d} \Big\{\int_{y_4}^{y_1(d)} F^P(d_A(y),y) f(y;\varepsilon)\,dy + \int_{y_1(d)}^{y_2(d)} F^P(d,y) f(y;\varepsilon)\,dy$$
$$+ \int_{y_2(d)}^{\infty} F^P(d,y) f(y;\varepsilon)\,dy\Big\} = 0.$$

Here $y_1(d) = d_A^{-1}(d)$, $y_2(d) = d_P^{-1}(d)$. When the differentiation is carried out, derivatives with respect to the limits cancel, and we get:

$$(2.43) \qquad \int_{y_1(d)}^{y_2(d)} \frac{\partial}{\partial d} F^P(d,y)f(y;\varepsilon)dy + \int_{y_2(d)}^{\infty} \frac{\partial}{\partial d} F^P(d,y)f(y;\varepsilon)dy = 0.$$

Differentiate (2.43) totally w.r.t. ε. By using the second order condition for a maximum, we get equation (2.41) at $\varepsilon = 0$.

Q.E.D.

The theorem is, of course, only a crude statement of necessary and sufficient conditions. The point is that the agent's preference function only enters the conditions via y_1, y_2, y_3 and y_4. If we look at our picture, (2.41) simply says that the loss that occurs to the principal when he shifts $\bar{d}$ up, comes from the region (y_1,y_2), and this loss has to be less than the gains that occur in the region (y_2,∞).

From the proof it should be clear how to state a necessary and sufficient condition for more general cases.

From Theorem 2.7 we infer that many factors determine how the optimal interval control changes, and so it is not easy to give very general conditions under which (2.41) and (2.42) hold. Looking only at (2.41), we see that if $\frac{\partial^2}{\partial d \partial y} F^P(\bar{d},y) > 0 \ \forall y$ (which is quite reasonable to assume with an increasing response function), then a mean-preserving spread of the distribution of $\tilde{y}$ will result in an increase in $\bar{d}$, when the mean of $\tilde{y}$ is below y_1.[20] Whether or not the mean is below y_1 depends on the closeness of the agent's and the principal's response functions. (Even if the mean is above y_1, a mean-preserving spread may result in an increased $\bar{d}$.) In Weitzman's model of economic control the response functions of the producer and the cneter intersect at the best centralized act d* (see (2.35), (2.36), noticing that the price is set at p = C'). Thus y_1 will always lie above the mean of $\tilde{y}$ (in Weitzman's model $\tilde{y} = \tilde{z}$). Consequently, the result that economic control will become less tight with an increase in the information gap does not depend on the assumption that $\tilde{y}$ is Normally distributed. Any mean-preserving spread in the information gap will result in increased freedom for economic agents.

Looking at other examples, it appears that responses must differ significantly around the mean value of $\tilde{y}$ in order for y_1 to fall below this mean. We infer then that it is quite often that increased uncertainty for the principal leads to increased freedom for the agent.[21]

2.3.6 Comparison of Preference Structure

The other component which determines the amount of delegation is the difference in objectives. Suppose we have two agents with the same information, but the other agent's preferences are in some sense closer to the principal's. Does this imply he will be given more freedom? The answer is in the affirmative, if we assume preference functions are unimodal. The appropriate notion of closeness is given by the following definition, which will provide a partial ordering of preference functions in terms of the principal's welfare.

Definition: Let A and A' be two agents with the same information $\tilde{y}$. We say that agent A:s preferences are uniformly closer than agent A':s w.r.t. the principal's, if for every $d \varepsilon D_0$

$$Y_{A'}(d) \subseteq Y_A(d) \subseteq Y_P(d),$$

or

$$Y_{A'}(d) \supseteq Y_A(d) \supseteq Y_P(d),$$

where

$$D_0 = \{d \varepsilon D | d_P(y) = d, \quad \text{for some } y\}.$$

We define uniform closeness only over the set of decisions which the principal may take (D_0), because the optimal interval controls will always be subsets of D_0.

Note that pairwise coherence is implied by the definition of uniform closeness.

The following lemma will clarify the meaning of uniform closeness.

Lemma 2.8: Suppose agent A:s preferences are uniformly closer than A':s w.r.t. the principal's, and that all preference functions are strictly unimodal. Then, for any interval control C,

$$d_{A'}(y|C) \leq d_A(y|C) \leq d_P(y),$$

or for every $y \varepsilon Y$.

$$d_{A'}(y|C) \geq d_A(y|C) \geq d_P(y),$$

Proof: Make the contrapositive assumption, say,

$$d_A(y|C) < d_{A'}(y|C) \leq d_P(y) \tag{2.44}$$

for some y. The other possibility is symmetric. Let $d' = d_{A'}(y|C)$. By (2.44) and unimodality, $y \varepsilon Y_P(d')$, $y \not\varepsilon Y_A(d')$. This implies, by uniform closeness, $y \not\varepsilon Y_{A'}(d')$, so $d_{A'}(y) < d'$. Since C is an interval, it contains all points between $d_A(y|C)$ and d'. By strict unimodality of agent A':s preferences he should then choose $d_{A'}(y|C) < d'$, since $d_{A'}(y) < d'$, which contradicts the definition of d'.

Q.E.D.

Notice that we have to restrict attention to interval controls. For controls with gaps the claim would be false in general. In that case we would need more information about the specific preference structure of the agents. With interval controls the relevant information is carried in the uncontrolled response functions alone.

The lemma shows that uniform closensss implies that agent A:s response function always lies <u>between</u> P:s and A':s. An example of uniform closeness would be the case where all three parties have exponential utility functions and A:s risk aversion coefficient lies between P:s and A':s. A picture of uniform closeness is given below. In this picture A is uniformly closer than A'.

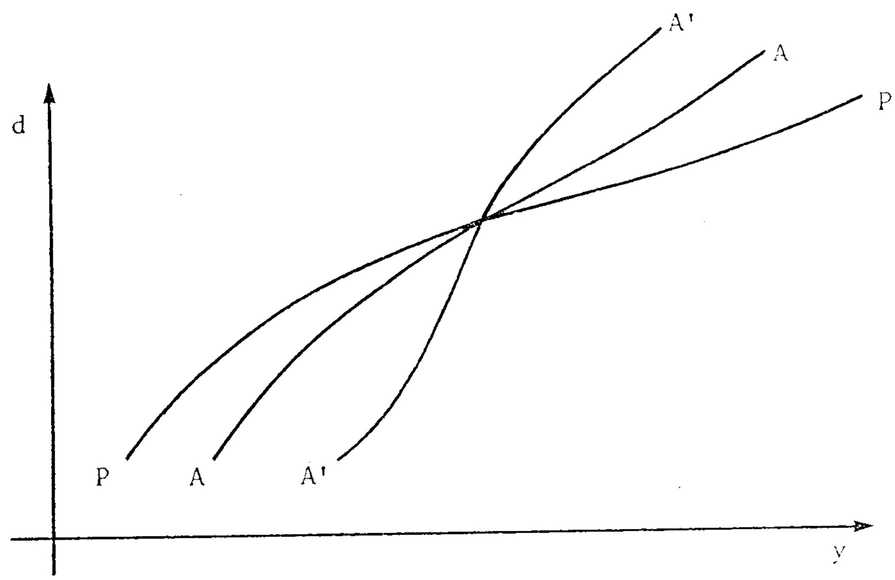

It should be noted that it is possible that A is uniformly closer in preferences than A' for one information signal $\tilde{y}$, but not for another. Hence, we cannot define uniform closeness directly over z.

The following is the main theorem on comparison of preferences.

Theorem 2.9: Assume preferences are concave and continuous. If agent A is uniformly closer than A' in preferences w.r.t. the principal, then the principal will be no worse off with agent A than with A', and he will give A at least as much freedom as A', regardless of the distribution of $\tilde{y}$.

Proof: The claim that the principal is no worse off with A than A' follows directly from Lemma 2.8, since A:s controlled response function is pointwise closer than A':s regardless of the control C.

For the second claim we need to show that if [a',b'] is an optimal control for A', then there exists an optimal control [a,b] for A, which contains [a',b']. By Lemma 2.6 the optimal intervals overlap, since they contain the principal's best centralized act. Hence, if [a,b] does not contain [a',b'], then either $a' < a$ or $b' > b$.

Assume $b' > b$. We will show that [a',b] is a strictly better control of A' than [a',b'], contradicting the optimality of [a',b']. Two cases are possible by uniform closeness.

Case 1: $Y_{A'}(b) \supseteq Y_A(b) \supseteq Y_P(b)$. An illustration of the situation is given below, where we have written $Y_1 = Y_{A'}(b)$, $Y_2 = Y_{A'}(b) - Y_A(b)$, $Y_3 = Y_A(b)$.

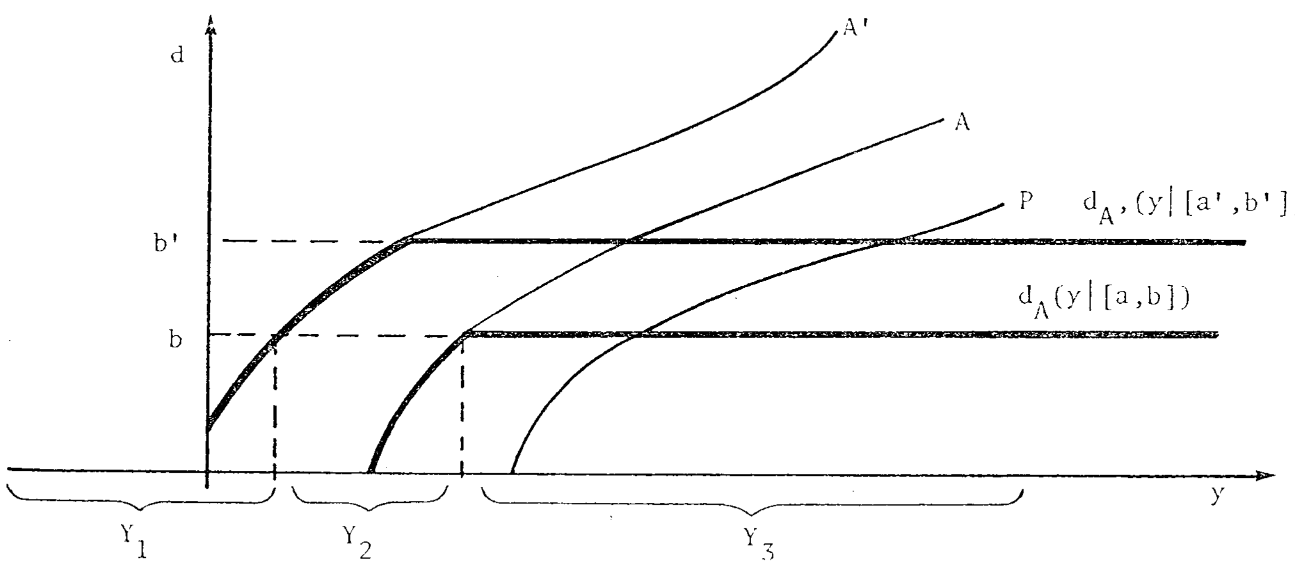

Recall that we write with lower case letters the expected preference functions; e.g., $f_P(d,y) = E(F_P(d,z)|y)$. Since F_P is concave so is f_P, and correspondingly for f_A and $f_{A'}$.

$$\int f_P(d_{A'}(y|[a',b']),y)dP(y) = \int_{Y_1} f_P(d_{A'}(y|[a',b']),y)dP(y)$$

(2.45)

$$+ \int_{Y_2} f_P(d_{A'}(y|[a',b']),y)dP(y) + \int_{Y_3} f_P(d_{A'}(y|[a',b']),y)dP(y).$$

By definition of $Y_{A'}(b)$,

$$(2.46) \quad \int_{Y_3} f_P(d_{A'}(y|[a',b']),y)dP(y) = \int_{Y_3} f_P(d_{A'}(y|[a',b]),y)dP(y).$$

On $Y_{A'}(b) - Y_A(b)$, $d_{A'}(y|[a',b']) \geq b$, whereas $d_A(y) \leq b$. Consequently, by uniform closeness and using Lemma 2.8, $d_P(y) \leq d_A(y) \leq b$. Since the principal's preference function is unimodal, he would prefer $d = b$ to $d_{A'}(y|[a',b'])$ on $Y_{A'}(b) - Y_A(b)$. But on this set, $d_{A'}(y|[a',b]) = b$, so

$$(2.47) \quad \int_{Y_2} f_P(d_{A'}(y|[a',b']),y)dP(y) \leq \int_{Y_2} f_P(d_{A'}(y|[a',b]),y)dP(y).$$

On $Y_A(b)$, $d_{A'}(y|[a',b'])$ is no better than $d_A(y|[a,b'])$ for the principal, by Lemma 2.8. On the other hand,

$$\int_{Y_3} f_P(d_A(y|[a,b']),y)dP(y) < \int_{Y_3} f_P(d_A(y|[a,b]),y)dP(y)$$

since $[a,b']$ was assumed suboptimal, and $d_A(y|[a,b']) = d_A(y|@a,b])$

on Y_3^c, Consequently,

$$(2.48)\quad \int_{Y_3} f_p(d_{A'}(y|[a',b']),y)dP(y) < \int_{Y_3} f_p(d_{A'}(y|[a',b]),y)dP(y).$$

Combining (2.45)-(2.48),

$$\int f_p(d_{A'}(y|[a',b']),y)dP(y) < \int f_p(d_{A'}(y|[a',b]),y)dP(y).$$

This contradicts the optimality of [a',b']. Hence, $b' \leq b$ in Case 1.

Case 2: $Y_{A'}(b) \subseteq Y_A(b) \subseteq Y_p(b)$. Proceeding analogously to Case 1, we can show that $b' > b$ leads to a contradiction of the optimality of [a',b']. Hence, $b' \leq b$ also in this case.

Finally, a similar argument shows that $a' \geq a$, completing the proof. Q.E.D.

If we look back at Example 2.2 we find that when two agents, A and A', have weights $b_{A'} > b_A \geq b_p$, then A:s preferences are uniformly closer than A':s, and so by the theorem above he is given more freedom. Similarly, the theorem applies to changes in B" and C" in Weitzman's model. When B" increases or C" decreases, the production unit's and the center's response functions get further apart and the control is tightened, verifying our earlier conclusions. This result is independent of the distribution of $\tilde{y}$.[22]

If we call two agents similar provided they always take acts on the same side of the principal's response function, we have the following partial converse of Theorem 2.9.

Theorem 2.10: Assume agents are similar and the agents' and the principal's preferences are concave and pairwise coherent. If the agents cannot be ordered by uniform closeness, then there exist two distributions of $\tilde{y}$, such that for one, agent A is strictly preferred to and given as much freedom as A', and for the other, the reverse holds true.

Proof: We will only outline the proof for increasing response functions. The coherence condition takes care of the general case.

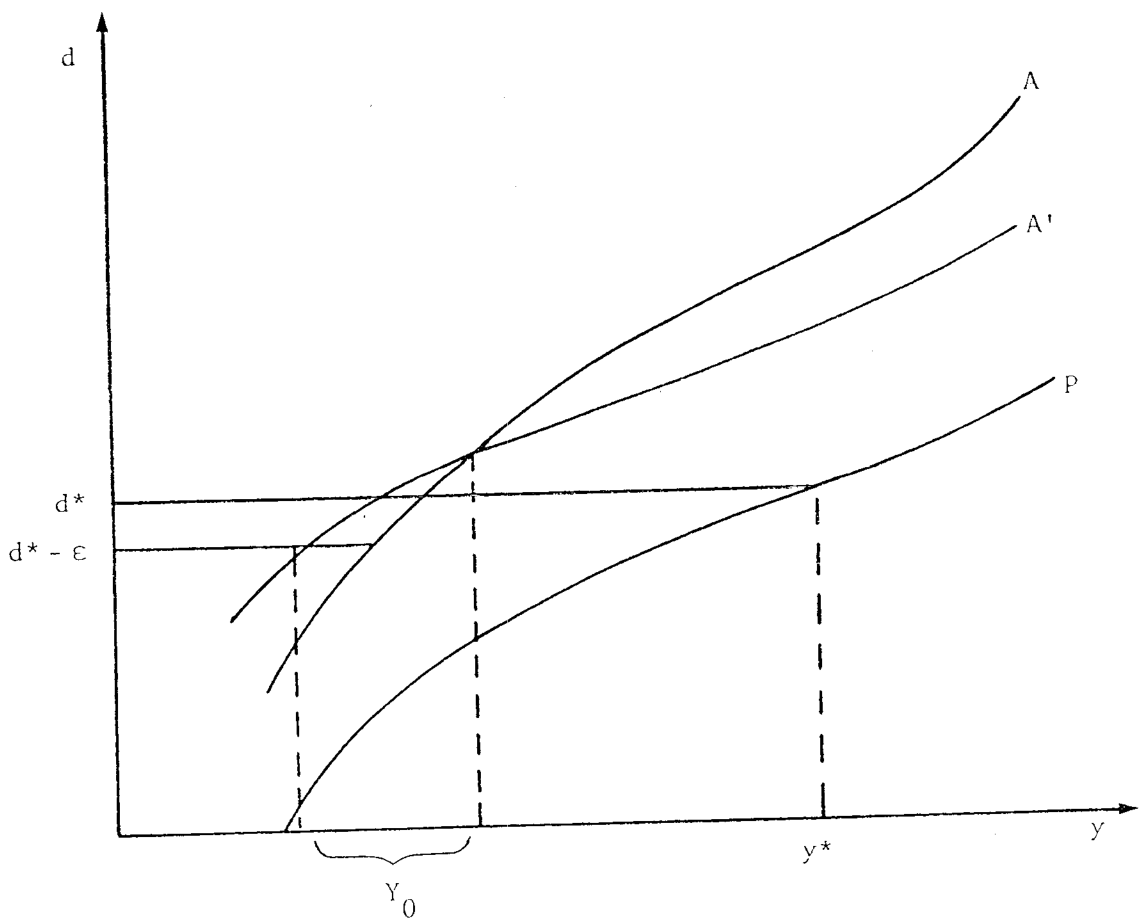

By assumption there exists a region Y_0 such that A is uniformly closer (strictly) than A' in Y_0. Make the distribution of $\tilde{y}$ such that d*, the principal's best centralized act will intersect $d_A(y)$ in Y_0; (we can assume $d_A(y)$ is not constant on Y_0 by increasing Y_0 if necessary and involving the similarity and coherence condition). This can be done by putting essentially all weight on $y^* = d_P^{-1}(d^*)$. However, let Y_0 have a positive probability measure, too. Then it pays to decentralize to some extent; in the picture above by using $[d^* - \varepsilon, d^*]$ as a control. (For an extended discussion, see the next section.) By coherence, A will have a uniformly closer response than A' on $Y_0 \cup \{y^*\}$, whatever control is used. On Y_0 it will be strictly closer. Consequently, A is given at least as much freedom as A' by Theorem 2.9, and he will be strictly preferred to A'.

Likewise, we can find a region Y_0', such that A' is uniformly closer than A in Y_0', and repeat the argument above.

Q.E.D.

As a consequence of Theorems 2.9 and 2.10 we get:

<u>Corollary 2.11</u>: Assume preferences are concave and pairwise coherent. Among similar agents, A is preferred to A' for all distributions of $\tilde{y}$ if and only if A is uniformly closer than A' in preferences w.r.t. the principal.

2.3.7 Decentralization vs. Centralization

When are there gains to decentralization? This is a basic question of interest. We will provide one simple sufficient condition for when it pays to decentralize using interval controls. Essentially, it says that if the agent takes acts both below and above the principal's best centralized act and preferences are coherent, then the principal should give the agent some freedom. Later, in Section 2.4, we will discuss how this condition can be weakened when transfer payments are allowed.

The following lemma, though obvious, is the main principle behind decentralization.

Lemma 2.12: Let d^* be the best centralized act. Assume there exists a control $C \varepsilon \mathcal{C}$, containing d^*, and a set $Y_0 \subseteq d_A^{-1}(C)$ for which:

(i) the principal prefers weakly $d_A(y|C)$ to d^* for every y;

(ii) the principal prefers strictly $d_A(y)$ to d^* for $y \varepsilon Y_0$;

(iii) $P(Y_0) > 0$.

Then the control C is strictly preferred to d^* both by the agent and the principal.

Proof: Obvious.

The point is that the agent's decision reveals information about $\tilde{y}$ (sometimes perfect information) and the lemma says that if the principal, conditional on the agent's decision, prefers it to the best centralized decision d*, then the agent should be allowed to make the decision. Notice that if there would be a set of decisions among which the principal is indifferent, a Pareto improvement would not be guaranteed by delegating this set to the agent (as is the case in the world with symmetric information). Conditional on the agent's choice, the indifference would generally change; possibly so that the agent's choice always becomes worse than a particular best centralized decision. This phenomenon is alien to adverse selection.

Theorem 2.13:[23] Assume:

(i) the expected preference functions f_P and f_A are unimodal, coherent and continuous;

(ii) the distribution of $\tilde{y}$ is nonatomic;

(iii) the agent takes acts both below and above the best centralized act d*;

(iv) $d_A(y) \neq d_P(y)$ for $y \varepsilon \{y | d_P(y) = d^*\}$.

Then there exists an interval control C such that both principal and agent are strictly better off with C than with d*; i.e., it is Pareto improving to decentralize.

Proof: The picture below indicates the simple idea behind the proof.

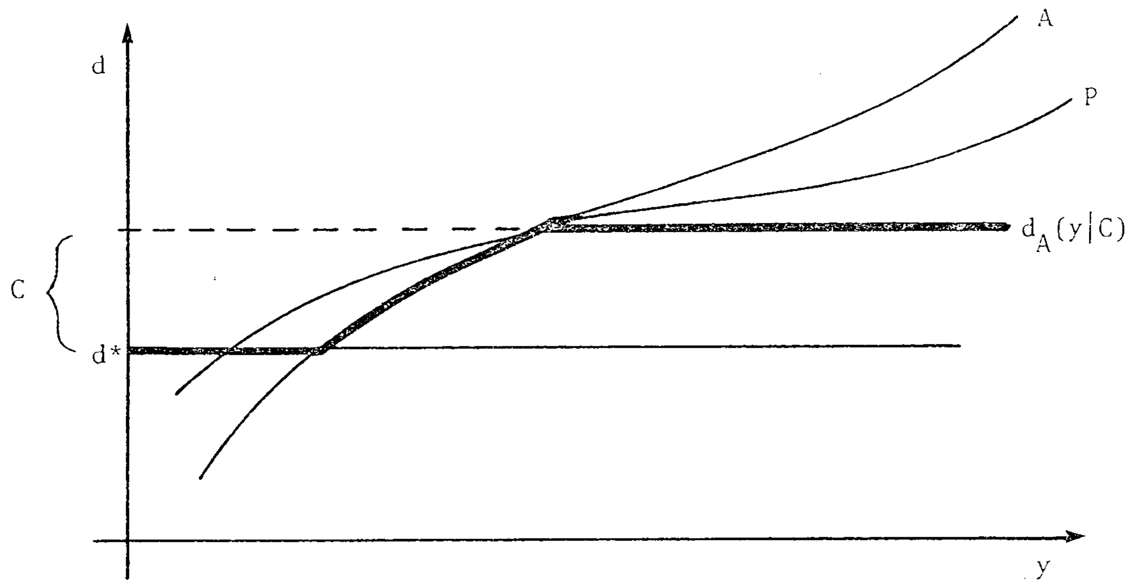

Clearly the heavy line dominates d = d* from the principal's and the agent's point of view.

We will show that we can find C and Y_0 as required in Lemma 2.12, with C an interval. Let $Y(\varepsilon) = \{y | d^* - \varepsilon \leq d_A(y) \leq d^* + \varepsilon\}$. By continuity, coherence and (iv), $\varepsilon > 0$ can be chosen small enough so that either $d_P(y) > d_A(y)$ or $d_P(y) < d_A(y)$ on $Y(\varepsilon)$.

Take the first case, $d_P(y) > d_A(y)$ on $Y(\varepsilon)$, as the second case is symmetric. Let $Y^+(\varepsilon) = \{y \varepsilon Y(\varepsilon) | d_A(y) \geq d^*\}$. By (iii) and continuity, $Y^+(\varepsilon)$ is of full dimension and so by (ii) $P(Y^+(\varepsilon)) > 0$. By strict unimodality the principal prefers strictly $d_A(y)$ to d* on $A^+(\varepsilon)$. Take $d_1 = d^*$, $d_2 = \sup \{d_A(y) | y \varepsilon Y^+(\varepsilon)\}$, and define $C = [d_1, d_2]$. Let y by arbitrary. Three possibilities arise:

(1) $y \varepsilon Y^+(\varepsilon)$. Then $d_A(y)$ is strictly preferred to d* by the principal.

(2) $d_A(y|C) = d^*$

(3) $d_A(y|C) = d_2$.

By coherence the third case implies $d_P(y) > d_2$ and so by unimodality the principal prefers d_2 to d^*. Hence $Y^+(\varepsilon)$ and C satisfy the conditions of Lemma 2.12. Q.E.D.

Let us emphasize that coherence is a crucial assumption in the theorem, as we already argued in connection with the definition. Indeed, if we say that the principal and the agent are noncoherent when $Y_A^C(d) \subseteq Y_P(d)$ or $Y_P^C(d) \subseteq Y_A(d)$ for every d, then the agent is of no value for the principal (see figure following the definition of coherence). This follows from the observation that if C is an arbitrary control, then the principal prefers any constant decision in C to the use of the control C.

It should be clear from Lemma 2.12 that condition (ii) could be replaced by a number of variants including some, which apply for discrete distributions.

We can get a weak converse of the theorem. Since it is trivially true that a necessary condition for decentralization by intervals is that the ranges of the principal's and the agent's response functions are not disjoint, we have:

Theorem 2.14: The agent is of value (i.e., it pays to decentralize) for all nonatomic distributions of $\tilde{y}$, if and only if preferences are coherent and the range of his response function contains the range of the principal's.

As an application of the theorems, let us look at an example of "management by participation," where the agent is given some freedom in choosing his own objectives as a result of his superior information about the complexity of the task.

Example 2:7 The agent's wage is determined according to a goal-based scheme (see Keren [1972]) as follows:

$$s(\bar{x}) = \begin{cases} w + b\bar{x}, & \text{if} \quad x \geq \bar{x} \\ w \quad , & \text{if} \quad x < \bar{x}, \end{cases} \tag{2.49}$$

where s = total salary

w = flat wage

b = bonus constant > 0

$\bar{x}$ = goal

x = output

Assume the production function is $x(e,\tilde{z}) = x(e) \cdot \tilde{z}$, $x(0) = 0$; e is the agent's effort and $\tilde{z} \geq 0$ is the randomness factor. For simplicity assume $\tilde{z} = \tilde{y}$, i.e., the agent knows z, whereas the principal is uncertain about it.

The agent's utility function is

$$u^A(s,e) = s - V(e) \quad ;V(0) = 0,\ V'(0) = 0,$$

where $V(e)$ measures disutility and is convex with a vertical asymptote at $\bar{e}$. The principal's utility function is linear over money without other attributes.

In our previous terminology the action variables are $\bar{x}$ and e (w is kept constant), and the control variable is $\bar{x}$ alone since e is assumed nonobservable. The derived preference functions are

$$
\begin{aligned}
F^A(\bar{x},e,y) &= w + b\bar{x} - V(e), \quad && \text{if } y \geq \frac{\bar{x}}{x(e)}, \\
&\quad w - V(e) \quad , && \text{if } y < \frac{\bar{x}}{x(e)}, \\
F^P(\bar{x},e,y) &= x - w - b\bar{x} \quad , && \text{if } y \geq \frac{\bar{x}}{x(e)}, \\
&\quad x - w \quad , && \text{if } y < \frac{\bar{x}}{x(e)}.
\end{aligned}
\tag{2.50}
$$

Let us look at response functions of $\bar{x}$ (i.e., what goal each player would prefer given y). We have:

$$\bar{x}_A(y) = x(e^*(y)) \cdot y,$$

$$\bar{x}_P(y) = x(\hat{e}(y)) \cdot y,$$

where $e^*(y)$ is the solution of $b \cdot x'(e) \cdot y - V'(e) = 0$, and $\hat{e}(y)$ solves $b \cdot x(e) \cdot y - V(e) = 0$.

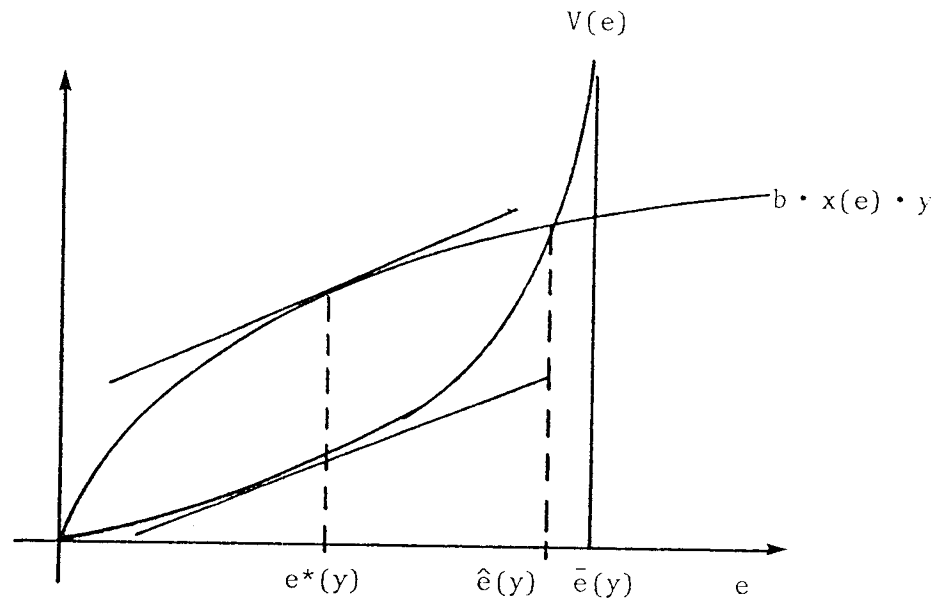

The principal always prefers a higher goal than the agent. By our assumptions $\bar{x}_P(0) = \bar{x}_A(0)$ and $|\bar{x}_P(y) - \bar{x}_A(y)| \to 0$ as $y \to \infty$. This implies that the response functions look like:

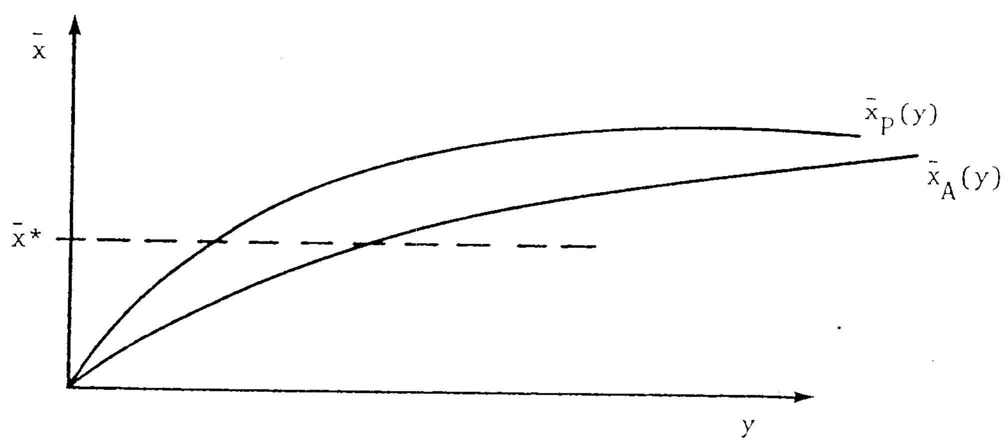

For any nonatomic distribution with range $[0,\infty)$, the agent will be given some freedom in deciding his goal, since it is clear that F_A and F_p, written as functions of $\bar{x}$ and y alone by substituting $e(\bar{x})$, are unimodal in $\bar{x}$. []

Going back to Weitzman's model, we see that Theorem 2.13 could have been used to prove Proposition 2.3. Equation (2.34) implies coherence, convexity was assumed, and the price p can be chosen so that (iv) holds. This shows how Theorem 2.13 can be applied even in a multi-dimensional control problem (both price and quantity are controllable in Weitzman's model), by fixing all but one of the control variables.

2.3.8 Motivation and Learning Aspects in the Use of Quantity Controls

It is often held that one of the major advantages of decentralization is the motivation it provides for the subordinate to do a better job. For instance, in a multi-division organization it is well-known that prices and quantities are equally good decision modes when the center has full information (under standard convexity assumptions), but prices are generally thought of as superior on motivational grounds. Exactly how this motivation works is rarely analyzed, rather it is taken as a matter of fact. On the other hand, looking at a model with fixed information structure, it is hard to find support for this view. One has to rely on arguments outside the standard economic framework to explain how motivation could improve the outcome.[24]

We will here show that if one views information acquisition as a main function of managers who participate in decision-making, then an aspect of motivation enters into the decentralization decision.

A manager performs two tasks. He makes decisions and supervises their implementation. Both activities take time and effort. We are here interested in the effort he spends on decision-making. The purpose of this effort expenditure is to improve his information about the decision problem he faces.[25] The amount of effort he will spend on this activity depends on how much freedom he has in making a decision, that is, how much responsibility he has been delegated. If he has no freedom and only is asked to implement a decision made by somebody else, he has no direct incentives to collect information about the problem (at least for decision-making purposes). It is only when he is delegated some decision-making power himself, that he will engage in acquisition of information.

Is it good or bad for the superior to have his manager become more informed? We looked at this question earlier, but could not conclude generally that a more informed manager was preferable. Yet, it seems natural to believe that this is the common case, and we will confine ourselves to situations in which this is true. For instance, in the production problem studied earlier (Example 2.2) we found that the principal was better off when the agent got more informed.

It is then clear that if the manager will spend more effort on collecting information when he is delegated more decision-making responsibility, the optimal amount of delegation will take this effect into account. In particular, he will be given more freedom in order to be motivated to acquire information. That the manager indeed will be motivated to get more informed will be shown below.

For simplicity we will assume that the manager's utility function is separable,

$$F^A(d,e,z) = F^A(d,z) - V(e),$$

where V(e) is disutility from effort. Let $\tilde{y}_1$ and $(\tilde{y}_1,\tilde{y}_2)$ be two information systems available to the manager at effort levels e_1 and e_2; $e_2 > e_1$, since the latter information system is strictly finer. Let C and C' be two interval controls with $C \subset C'$. Suppose the manager prefers $(\tilde{y}_1,\tilde{y}_2)$ to $\tilde{y}_1$ under the control C, so that

$$(2.51)\quad E[E[F^A(d(y_1,y_2|C),z) - F^A(d(y_1|C),z)]|y_1,y_2] > V(e_2) - V(e_1).$$

When C is enlarged to C', it follows by concavity of the preference function that:

$$(2.52)\quad d(y_1,y_2|C') - d(y_1|C') \geq d(y_1,y_2|C) - d(y_1|C), \qquad \forall (y_1,y_2).$$

Since $d(y_1,y_2|C)$ $(d(y_1,y_2|C'))$ maximizes the integral in (2.51) (i.e., F^A conditional on (y_1,y_2)) subject to the constraint C(C'), (2.52) implies that

$$(2.53)\quad \begin{aligned} &E[F^A(d(y_1,y_2|C'),z) - F^A(d(y_1|C'),z)|y_1,y_2] \geq \\ &E[F^A(d(y_1,y_2|C),z) - F^A(d(y_1|C'),z)|y_1,y_2]. \qquad \forall (y_1,y_2). \end{aligned}$$

Taking the expectation of (2.53) and using (2.51) we find that also under the constraint C', the manager prefers $(\tilde{y}_1,\tilde{y}_2)$ to $\tilde{y}_1$.

Assuming increased levels of effort result in increasingly finer information systems, the argument above shows that with more freedom the manager will invest at least as much effort into information acquisition as before. Generally he will invest more and get a strictly finer information system, which we assumed is to the benefit of the principal.

Consequently, in determining the optimal level of delegation, the principal should also make a provision for this motivational aspect of the problem, and give the agent more freedom than he would if the agent's information were independent of effort.[26] In terms of price vs. quantity controls we conclude that prices have a comparative advantage as motivational tools as is often claimed.

Another aspect which may increase the amount of freedom is learning. We have formulated our problem in a static framework, but suppose there were two decision periods instead of one. If the agent's characteristics (his expertise and preferences) were perfectly known to the principal, nothing would be essentially changed and our previous analysis would apply to determine the optimal level of delegation.[27] However, if the principal is not fully informed about his agent, the outcome of the first period will be a signal about the agent's characteristics, and the principal would have to take this into account when determining the optimal control set in the first period. Assuming the agent optimizes his problem period by period without thinking about strategic behavior, it is clear that the principal will give the agent at least as much freedom in the first period, as he would have done in a one-period problem. More freedom will provide finer information about the agent and this can never hurt the principal by Blackwell's theorem. Since learning in dynamic models has been explored in several papers,[28] and the point is rather obvious anyhow, we will not analyze the issue further.

To summarize the discussion, we have identified three components of the optimal interval control, schematically described in the picture below.

| use of agent's superior information | motivation | learning |

First, the principal delegates decision-making responsibility in order to utilize the agent's superior information. Generally it will be restricted because of differences in objectives. Secondly, the principal may expand delegation to motivate the agent to acquire more information. And thirdly, the principal can use delegation for learning about the agent's characteristics.

2.4 Controls with Compensation

In the one-dimensional control problem, the thrust was on studying optimal interval controls and their characteristics as a function of the difference in information and preferences. Extensions to multi-dimensional control problems are possible, but much weaker results will be obtained. This motivates us to look at a particular case of a two-dimensional control problem, namely one with $d = (a,w)$, where a is the action variable of the underlying problem structure and w is a compensation variable. We will call this the regulation problem following Weitzman [1976b]. It has the same structure as screening and signalling models, but we will look at centralized solutions rather than market solutions.

2.4.1 The Regulation Problem

As we already saw in Example 2.3, the regulation problem reduces to finding an optimal control function $s(a)$. Let the preference functions be $F^A(a,w,z)$ and $F^P(a,w,z)$. We assume $F_2^A > 0$, $F_2^P \leq 0$, $F_{11}^A < 0$, $F_{11}^P < 0$. The regulation problem can then be written:

$$(2.54) \qquad \max_{a(y),\, s(y)} E[F^P(a(y),\ s(a(y)),z)]$$
$$\text{s.g.} \quad a(y) = \underset{a}{\operatorname{argmax}}\, E[F^A(a,s(a),z)\,|\,y], \qquad \forall y,$$

assuming the constraint is well-defined for almost every y.

This problem appears generally very difficult to solve. Two particular difficulties can be seen directly. First, the constraint involves a global solution to an optimization problem, and even with concavity assumptions on F^A, we cannot replace the constraint with the first-order condition:

$$(2.55) \quad f_1^A(a(y),s(a(y)),y) + f_2^A(a(y),s(a(y)),y) \cdot s'(a(y)) = 0, \ \forall y,$$

since s may make the objective function of the agent nonconcave in a. An additional assumption, used in screening models, will resolve this problem, as we will see shortly. The other problem is that the control function s is a function of y via a(y), which is unknown. One way of approaching this problem, is to ignore it at first and write s(a(y)) = g(y), letting g(y) vary freely. If the optimal a(y) is strictly monotone, we can always find the s that corresponds to the optimal g(y) from the relationship above. However, if a(y) is not monotone we have exceeded our degrees of freedom.

Under certain simplifying assumptions even the second problem can be solved; for instance, if we assume that the agent's utility function is additively separable and linear in w (see Spence [1977]). But generally the difficulty remains and we will therefore turn to another solution approach, which carefully considers the agent's global maximization and yields a partial solution to the regulation problem under some further assumptions.[29] What we will do is to study what response functions a(y) the agent can be induced to take with different control functions s; phrased differently, we will look at the attainable set of controlled response functions. The work is based on Riley [1976], and it will be applied to generalize some of the results in Weitzman [1976b].

Look again at the first order condition for the agent's maximization problem; Equation (2.55). To find out a'(y), differentiate (2.55) totally w.r.t. y which gives:

$$\frac{\partial}{\partial a}(f_1^A + f_2^A \cdot s') \cdot a' + \frac{\partial}{\partial y}(f_1^A + f_2^A \cdot s') = 0. \tag{2.56}$$

If (2.55) is to give a local maximum we must have

$$\frac{\partial}{\partial a}(f_1^A + f_2^A \cdot s') = \frac{-1}{a'} \cdot \frac{\partial}{\partial y}(f_1^A + f_2^A \cdot s') < 0. \tag{2.57}$$

We have

$$\frac{\partial}{\partial y}(f_1^A + f_2^A \cdot s') = \frac{\partial}{\partial y}\left[\left(\frac{f_1^A}{f_2^A} + s'\right) \cdot f_2^A\right] = \frac{\partial}{\partial y}\left(\frac{f_1^A}{f_2^A}\right), \tag{2.58}$$

since $f_2^A > 0$ and using (2.55). Substitution in (2.57) gives

$$\frac{\partial}{\partial y}(f_1^A + f_2^A \cdot s') = \frac{-1}{a'} \cdot \frac{\partial}{\partial y}\left(\frac{f_1^A}{f_2^A}\right) < 0. \tag{2.59}$$

(2.59) indicates some of the restrictions that one faces when trying to control the agent. Assume, as is frequently done in the theory of screening, that $\frac{\partial}{\partial y}\left(\frac{f_1^A}{f_2^A}\right) > 0$. Then we can only induce the agent to take increasing response functions. On the other hand, we can prove the converse too.

Theorem 2.15: Assume $\frac{\partial}{\partial y}\left[\frac{f_1^A}{f_2^A}\right] > 0$. Then any differentiable nondecreasing response function is attainable. Conversely, any attainable response function must be nondecreasing.

Proof: We will prove the first part of the theorem assuming first that $a'(y) > 0$ for all y, and then argue that $a'(y) = 0$ can also be allowed. If $a'(y) > 0$, we can equivalently write (2.55) as

$$(2.60) \qquad f_1^A(a,s(a),y^{-1}(a)) + f_2^A(a,s(a),y^{-1}(a)) \cdot s'(a) = 0$$

$$\forall a \varepsilon \ \{a = a(y) \mid y \varepsilon Y\},$$

where Y is the range of y. This is a differential equation which has a unique solution s(a) under our assumptions. This means that for any a(y) with $a'(y) > 0$, there exists a control function s(a), which at least yields the first order condition (2.55). From (2.59) we know that it moreover gives a local maximum along a(y). To prove that each local maximum is, in fact, global, study the picture below:[30]

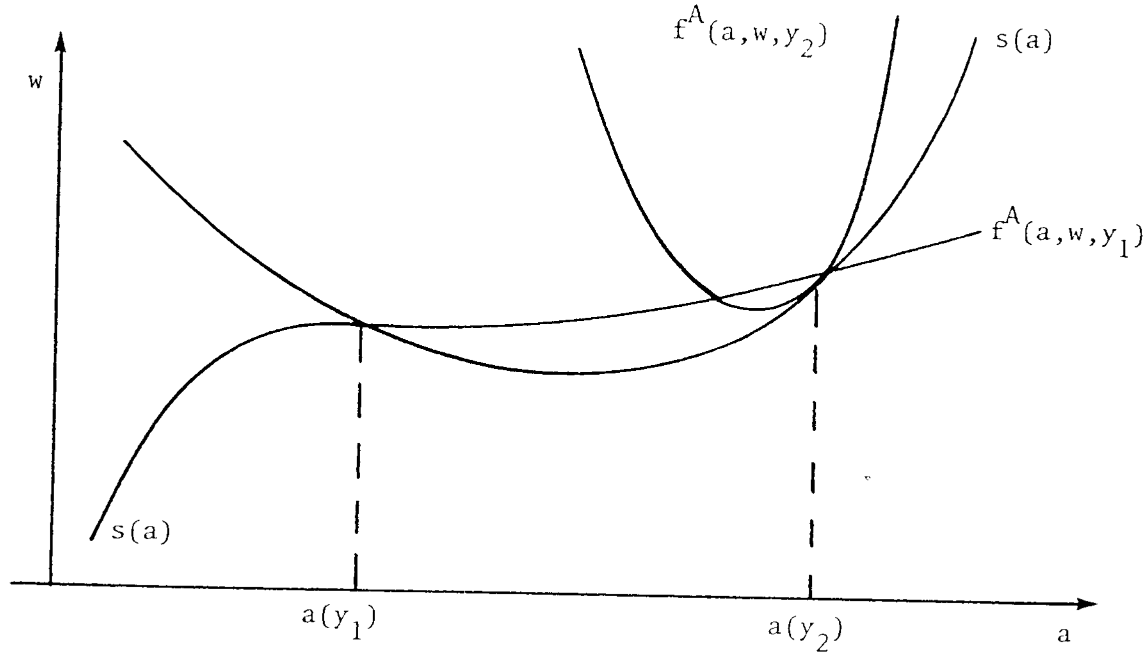

A contrapositive assumption would imply that the indifference curve of f^A, which is tangential at $a(y_1)$, intersects s(a) at some other point $a = a(y_2)$. By construction of s(a), there is a tangential indifference curve at $a(y_2)$, and this would imply $\frac{f_1^A(a(y_2),y_1)}{f_2^A(a(y_2),y_1)} > \frac{f_1^A(a(y_2),y_2)}{f_2^A(a(y_2),y_2)}$, (if $y_2 > y_1$; otherwise reversed) in contradiction to our assumption $\frac{\partial}{\partial y}\left(\frac{f_1^A}{f_2^A}\right) > 0$; since note that $-\frac{f_1^A}{f_2^A}$ is the slope of the indifference curves. From the geometry it is also clear that the agent can be kept at a constant a, i.e., $a'(y) = 0$ is possible, by decreasing s(a) sufficiently rapidly. This proves the first part of the theorem.

The slope interpretation indicates immediately that the agent will never take lower values of a when y increases. This proves the second part of the theorem.

Q.E.D.

The assumption $\frac{\partial}{\partial y}\left[\frac{f_1^A}{f_2^A}\right] > 0$ is natural in many situations. For instance, in the insurance model (Example 2.3) it says that higher-risk individuals have a lower opportunity cost for buying insurance; in Spence's model on job market signalling, the assumption is that individuals with higher ability have a lower opportunity cost for education.

Briefly, Theorem 2.15 says that the attainable set of controlled response functions equals the class of nondecreasing response functions. Of course, this is only a partial solution to the regulation problem in general, since we have not considered the costs associated with different controlled response functions. In some cases it yields a complete solution.

As an application, we look at the case where the preference functions can be written $f^A(a,y) + w$, $f^P(a,y) - w$; i.e., when there are no wealth effects and utilities are transferable in money. The (ex ante) Pareto optimal solution follows by solving the problem:

$$\max_{a(y),s(a)} \int \{f^P(a(y),y) + f^A(a(y),y)\}dP(y),$$

$$\text{s.t.} \qquad f_1^A(a(y),y) + s'(a(y)) = 0 \qquad \forall y.$$

(the weights of f^P and f^A have to equal 1 with transferable utilities). Assume $f_{12}^A > 0$, $f_{12}^P > 0$, which implies that the uncontrolled response functions are increasing. The first best solution $a^*(y)$ can be found by pointwise maximization of the integrand. It is easily seen to be increasing, since $f^A_{12} > 0$ and $f^P_{12} > 0$. Hence, Theorem 2.15 implies that there exists a control function s, which yields a*(y) as the agent's controlled response. The first best solution can be attained since s is just a transfer, and so there are no efficiency losses due to differential information. The statement applies, for instance, to the production problem (Example 2.2), if we make the assumption of transferable utilities.

Whenever the principal's preference function is independent of s (which was the case in the economic planning model, Section 2.3.3) and the agent satisfies $\frac{\partial}{\partial y}\left(\frac{f^A_1}{f^A_2}\right) > 0$, the constraint in the regulation problem can be replaced by the requirement that a(y) is nondecreasing. This is a substantial simplification and it is straightforward to see that the form of the optimal solution will look as in the picture below:

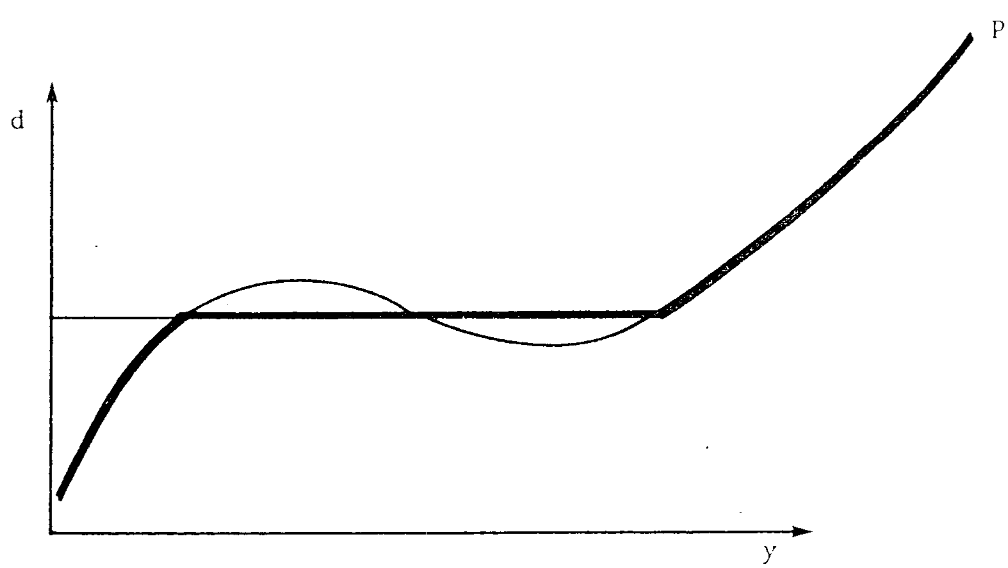

The optimal controlled response will follow the increasing part of the principal's response function, and take flat jumps over decreasing parts.

We can apply this discussion to an extension of Section 2.3.3 studied by Weitzman [1976b]. Suppose there are n producers instead of one. Each controls one decision d_i, and knows its own cost function $C(d_i;y_i)$. W.l.o.g. we can assume $C_{12} \leq 0$; i.e., y_i reflects decreases in marginal costs. y_i's are unknown to the center. The producers' decisions are interrelated through a joint benefit function B(d), $d = (d_1, \ldots, d_n)$. The center can announce price functions $p_i(d_i)$ with the objective to maximize the expected value of $B(d) - \sum_{i=1}^{n} C(d_i,y_i)$, given that divisions will maximize $p_i(d_i) - C(d_i,y_i)$. An upper bound for what the center can achieve is given by the optimal team solution without communication (see Marschak and Radner [1972]). This solution specifies optimal response functions $d_i^*(y_i)$ for the producers. It is quite natural that these functions are increasing (with our assumption $C_{12} \leq 0$). If y_i's are independent this is the case, but even with dependencies it is presumably rare that one producer should decrease the production as the marginal cost decreases. In any case we can apply Theorem 2.15 to this coordination problem. We conclude that the organization can be made to perform as well as a team without communication possibilities, if and only if the optimal team solution is nondecreasing. This is a generalization of part of Weitzman's results.

In Chapter III we will study Groves' scheme, which is designed to achieve the optimal team solution under essentially any communication structure. However, for Groves' scheme to work we have to assume that the agents' information signals are independent. If we make that assumption when no communication takes place, the solution becomes rather trivial. What is interesting is that there are situations with dependent signals for which the team solution can be achieved. We will return to this issue in Chapter III.

Usually it is hard to determine the optimal team solution.[31] By using quadratic approximations of benefit and cost functions (see (2.15)-(2.16)), and looking at a restricted class of joint probability distributions over signals (which includes the joint Normal distribution), Weitzman shows that the optimal response functions are linear. Provided they are increasing, one can induce firms to follow them using quadratic price schedules of the form:

$$(2.61) \qquad p_i \cdot d_i + w_i \cdot (d_i - d_i^*)^2, \qquad i = 1, \ldots, n,$$

where $d^* = (d_i^*, \ldots, d_n^*)$ is the best centralized decision. The w_i's are weights, which are determined from a set of simultaneous equations.

(2.61) represents another kind of mixed price-quantity control. If the weights are set very high we have essentially a quantity order, and if they are equal to zero we have a pure price control. The weights depend on the curvature of costs and benefits, in a way which merely confirms our previous conclusions. More interestingly, Weitzman shows that the weights increase with positive correlation between marginal costs of firms. The reason is that prices tend to lead to over-reactions. Whenever one firm is increasing production because its costs are low, so do the others. Compared with a situation of independent marginal costs, more stabilization is needed.

Weitzman's model is not well adapted to an analysis of the effects of changes in the information gap. This is best understood by looking at the case of independent signals. Then the quadratic price schedule can transmit the center's objective function completely, and this function is independent of how big the information gap is. There will be some second-order effects when signals are dependent, but they are not essential. The point is that changes in the information gap become important only when the center can transmit incompletely its information, which is the case, for instance, when it has to use interval controls.

2.4.2 Decentralization vs. Centralization

As in the one-dimensional control problem, we can again ask under what conditions the agent will be of value to the principal, i.e., when it pays to decentralize. With an enriched contractual space, the possibilities for successful decentralization are, of course, increased. In fact, they are increased rather substantially, since now the response functions need in no sense be close to each other. What counts is how the agent's rate of substitution between acts (a) and compensation (w) changes with y.

In Weitzman's model (Section 2.3.3) we say how we could manipulate the agent's response function using a price so that it intersected the principal's best centralized act and made decentralization valuable according to Theorem 2.13. The same technique can be used for an extension of the result to the following theorem.

Theorem 2.16: Assume $\frac{\partial}{\partial y}\left[\frac{f_1^A}{f_2^A}\right] > 0$. If the principal's response function is increasing, then it always pays to decentralize using a price scheme combined with an interval control.

Proof: Study the picture below. $\bar{a}$ is the principal's best centralized act, and $\bar{w}$ is the agent's wealth without compensation. Let $\bar{y} = d_p(\bar{a})$, i.e., the outcome of $\tilde{y}$ at which the principal would choose his best centralized act. $\bar{p}$ is the agent's MRS between acts

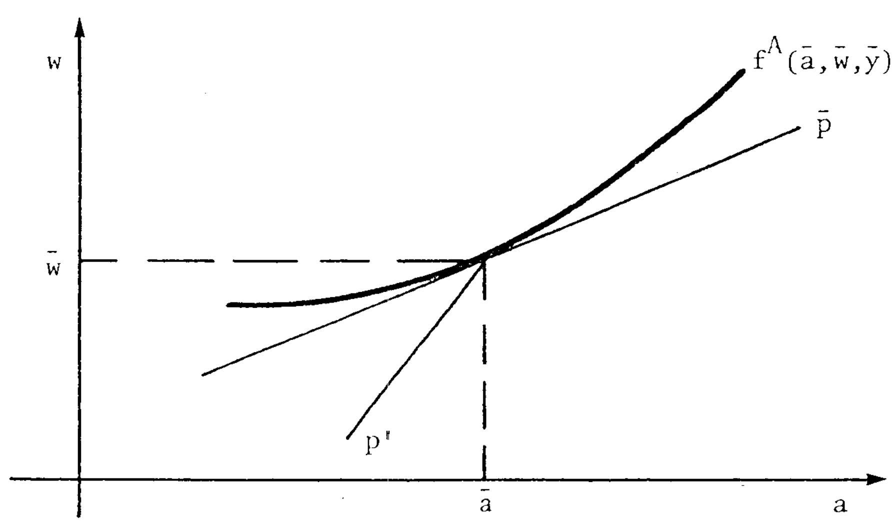

and wealth at $(\bar{a},\bar{w},\bar{y})$. If the agent would be faced with a price scheme $\bar{p} \cdot a - \bar{p} \cdot \bar{a}$, he would, by our assumption $\frac{\partial}{\partial y}\left[\frac{f_1^A}{f_2^A}\right] > 0$, have an increasing response function, for which $d_A(\bar{y}) = \bar{a}$. This may not yet make the principal better off (even if the agent is restricted to take acts in an interval around $\bar{a}$), since the agent's response function may be too steep when $\bar{p}$ is used. However, with a higher price p', the agent's response function can be made to intersect $a = \bar{a}$ to the left of $\bar{y}$ (see figure below). Using $p' \cdot a - p' \cdot \bar{a}$ as a price and $(\bar{a} - \varepsilon, \bar{a})$ (for some small $\varepsilon > 0$) as a decentralization interval, the principal will be better off. He will prefer the agent's response function to $\bar{a}$, since it is uniformly closer, and moreover he will receive payments from the agent when the agent takes acts different from $\bar{a}$. The agent will be no worse off with this arrangement, since he can always take $\bar{a}$ and pay nothing.

Q.E.D.

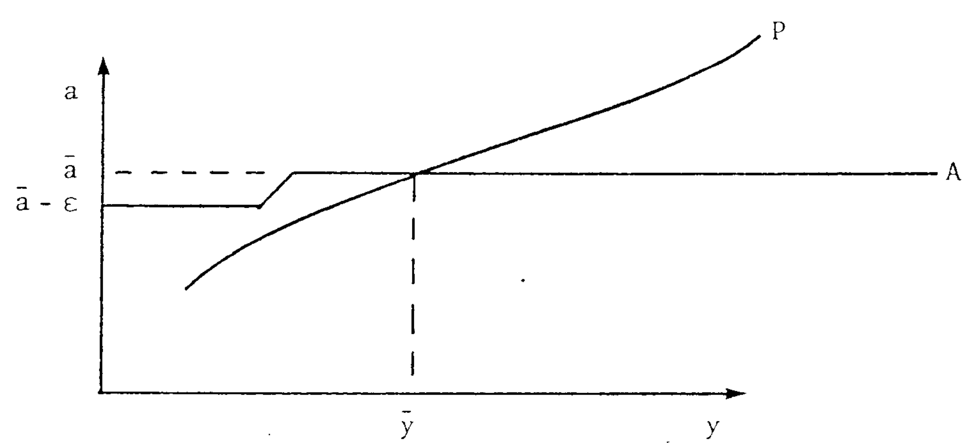

An obvious corollary of the theorem is that if $\frac{\partial}{\partial y}[\frac{f_1^A}{f_2^A}] < 0$ and the principal's response function is decreasing, then it pays to decentralize. The theorem could also have been proved under one of the two assumptions:

$$\text{(i)} \qquad \frac{\partial}{\partial y}[\frac{f_1^A}{f_2^A}] > 0 \qquad \text{and} \qquad \frac{\partial}{\partial y}[\frac{f_1^P}{f_2^P}] < 0,$$

or

$$\text{(ii)} \qquad \frac{\partial}{\partial y}[\frac{f_1^A}{f_2^A}] < 0 \qquad \text{and} \qquad \frac{\partial}{\partial y}[\frac{f_1^P}{f_2^P}] > 0,$$

since (i) implies that the principal's response function is increasing, and (ii) that it is decreasing when the transfer payment is kept constant. The role of these two conditions is very similar to the coherence condition we needed in the one-dimensional case. For instance, (i) says that when y increases, the agent is willing to take higher acts at lower compensation, whereas the principal is willing to pay more for higher acts.

Consequently, their interests move in the same direction with changes in y, and this provides an opportunity for Pareto improvements via decentralization.

There are many situations where either (i) or (ii) are natural assumptions. One example would be banks lending to firms which have superior information about the return of their investments (see Section 2.3.1). With higher expected returns, the firm would be willing to buy additional capital at a higher interest rate, at the same time as the bank would be willing to sell it at a lower interest rate. This is equivalent to condition (i).

Another example would be the model of management by participation, which we presented in Example 2.7. The rationale for letting the agent change his own target level is the fact that when returns from increased effort are higher, the agent is willing to raise the target (and consequently his effort level) for a lower bonus, whereas the principal is willing to pay a higher bonus for such a raise.

Notice, however, that neither (i) nor (ii) is generally satisfied in the screening models. For instance, in the insurance model (Example 2.3) this is not the case. Normally there will still be returns to decentralization (though there may not be). This does not imply that screening will be ineffective or not used in insurance markets. Our proposition applies in a centrally planned context where we are looking for Pareto improvements. In the market economy, screening is a consequence of firms looking for better alternatives, and generally some agents would prefer a centralized solution.

2.5 Efficiency with Differential Information

Little has so far been said about the efficiency of the solution to the delegation problem. The reason is that no unambiguously correct definition of efficiency is available under conditions of differential information. The traditional notions of efficiency are not normally operative when asymmetric information is present. For instance, the delegation solution is generally inefficient according to the standard of perfect information (i.e., the agent's information) and the same holds true for other decentralized solutions as well. And yet it seems that the delegation solution should be considered efficient in a constrained sense, because by definition it is the best that the principal can expect to attain with the means at his disposal.

In this section we will discuss the difficulties involved with a definition of efficiency under differential information. We will also propose a new notion of constrained efficiency, for which the delegation solution is efficient. The purpose of this definition is not so much to advocate that it is the appropriate notion of efficiency in general (since in our opinion it has some undesirable features), as it is to characterize those control sets in the delegation problem which cannot be changed for a joint improvement of the principal's and the agent's utilities after the agent has observed his signal. So far we have only studied the problem from the principal's point of view, and this analysis will indicate how to take account of the agent's interests as well.

Verbally stated, a decision (e.g., an allocation) should be considered efficient only if there is no other decision which everybody would prefer. In a world where an agent's private information is of no interest to the others (as is the case when there is no state uncertainty and the agent's private information is his preference ordering), the application of this definition is straightforward, since each agent can state his preference over decisions independently of the preferences of other agents.

Under differential information this is no longer the case. Now agents' preferences may reveal at least partly their private information, and consequently be of interest to others. In particular, when a person is willing to accept a change in the current decision, this will tell something about what he knows. Such implicit information transmission could be carried out, for instance, by market prices, and this has been recently studied under the notion of self-fulfilled expectations equilibria (e.g., Grossman [1976]).

A simple example will illustrate the point (found in Wilson [1977]). The economy consists of two agents, 1 and 2, whose endowments are contingent on the state of nature, which is either a or b. In state a agent 1 gets one dollar and agent 2 nothing; in state b the roles are reversed. We write this as $e = (e_a, e_b)$, $e_a = (1,0)$, $e_b = (0,1)$. Agent one knows the state of nature, agent two does not, and assesses equal probabilities to both states. The decision in this case is a sharing rule, which determines how the agents are going to split the dollar in each of the two states. The structure is compactly written in the following table:

	a	b	
1	1	0	{a},{b}
2	0	1	{a,b}
	1/2	1/2	

We are studying whether there is an alternative sharing rule $s = (s_a, s_b)$, $s_a = (1 - p, p)$, $s_b = (q, 1 - q)$, which both would prefer to the endowment e.

Suppose a has occurred. Then agent 1 demands p = 0 (cannot be negative), requiring q = 0 and so e cannot be dominated by any other sharing rule. This is in accordance with the traditional notion of efficiency. Suppose b occurs. Then e is dominated by s = ((.5,.5),(.5,.5)), if we follow traditional reasoning. Agent 1 prefers .5 to 0, and agent 2 (if risk-averse) prefers his certainty equivalent to the original lottery. But upon reflection agent 2 will consider s unacceptable. He should recognize that if agent 1 accepts s, it must be because state b has occurred. Given this inference, agent 2 should no longer accept s. Studying any other combination of p and q gives the same conclusion: agent 1 wants a change only in the case state b occurs, but this is the state in which agent 2 does not want any changes. There is no sharing rule s which both agent 1 and agent 2 would prefer simultaneously. In this sense we would be inclined to say that e is efficient.

The example suggests a definition of efficiency, where agents condition their acceptance on the information that is revealed by other agents' joint acceptance. The problem is that this logic leads to a noncooperative game of incomplete information. In order to know whether to accept or reject a new proposal, an agent has to form his beliefs about the other agents' acceptance behavior as a function of their private information. The natural formation of beliefs implies a set of indicator functions, one for each agent, which tells under what private information agents accept (respectively reject) the new proposal, and which moreover constitutes a Nash equilibrium in the sense that each agent's acceptance function is a best response against the others'.

Formally, let d be the current decision and d' a new proposal. Let $I_i(y_i;d,d')$, $i = 1, \ldots, n$, be a set of indicator functions and define $I^i(y^i;d,d') = \prod_{j\neq i} I_j(y_j;d,d')$, $I(y;d,d') = \prod_{i=1}^{n} I_i(y_i;d,d')$. Assume that these indicator functions form a Nash equilibrium in the sense that:

$$I_i(y_i;d,d') = 1, \quad \text{if and only if}$$

$$E(F_i(d',z)\,|\,y_i,\ I^i(y^i;d,d') = 1) \geq \tag{2.62}$$

$$E(F_i(d,z)\,|\,y_i,\ I^i(y^i;d,d') = 1), \qquad \forall y_i,\ \forall_i.$$

This condition embodies the idea that an agent should have no regrets when he accepts d' and finds out that the others did the same, as well as the idea that his acceptance behavior is determined by self-interest and is a best response against the other agents' acceptance strategies.

We could then define efficiency as follows:

Definition: A decision dεD is inferentially efficient at $\tilde{y} = y$, if there does not exist a decision d'εD and a set of indicator functions $\{I_i(y_i;d,d')\}$ as defined in (2.62), for which $I(y;d,d') = 1$, and for which (2.62) holds with strict inequality for at least one agent i.

It is clear that the table in the preceding example is inferentially efficient at both s = a and s = b. This property holds true more generally in cases with two agents, one of which has strictly superior information; namely, if d is efficient under perfect information for each outcome of $\tilde{y}$, then it is inferentially efficient at each $\tilde{y} = y$. The argument is the same as we used in the example.

We have argued that a logical extension of the traditional notion of efficiency leads to a cumbersome definition, which embodies a paradoxical element of noncooperative game theory.[32] Clearly, this cannot be considered very satisfactory. Moreover, one can show in the context of the delegation problem that a decision $d \varepsilon D$, which is inferentially efficient, can be dominated by a decision mechanism $N' = (d',Y)$ in the sense that both the principal and the agent can be made better off by employing N' instead of staying at d. This suggests that one should talk about decision mechanisms (or their outcome functions) rather than decisions as being efficient. Corresponding to our previous definition we would have:

Definition: A decision mechanism N = (d,Y) is inferentially efficient at $\tilde{y} = y$, if there does not exist another decision mechanism $N' = (d',Y)$ which everybody would prefer at $\tilde{y} = y$.[33]

For brevity we have not spelled out mathematically what is meant by everybody preferring N' to N. From the earlier definition this should be easily understood. Again, we look for a set of indicator functions which form a Nash equilibrium observing that an agent should take into account that the information changes both in the case N' is rejected or accepted. Thus the outcome of the game N will change when N' is rejected, from what it would have been, had N been played before voting for the change to N'. For this reason, an agent may be worse off simply by the proposal of N', even though he can always veto the implementation of N'.

Using an idea similar to the one employed by Wilson [1977] in his definition of conditional efficiency one can avoid such problems of inference employing the following definition:

Definition: A decision mechanism N = (d,Y) is weakly efficient, if there does not exist another decision mechanism N' = (d',Y), which everybody prefers in all states of their private information, i.e.,

$$E[F_i(d'(y),z)|y_i] \geq E[F_i(d(y),z)|y_i],$$

for every $\tilde{y}_i = y_i$, with strict inequality for some agent i, in some information state of $\tilde{y}$.

Here it is understood that d(y) and d'(y) are both decision and outcome functions of N and N'. If N' dominates N under this definition it is clear that no information will be revealed from accepting N'. However, there may be information states $\tilde{y} = y$ such that even though

N is weakly efficient it is not inferentially efficient. In this sense weak efficiency is a weaker requirement than inferential efficiency. In particular, if N is inferentially efficient at each $\tilde{y} = y$, it is, of course, also weakly efficient.

For the delegation problem (with the agent having strictly superior information) we can show that weak and inferential efficiency are equivalent notions. As we have argued earlier we can identify decision mechanisms with control sets in this case. We have:

Theorem 2.17: A control set C is weakly efficient if and only if it is inferentially efficient for each $\tilde{y} = y$.

Proof: The sufficiency is obvious. To prove necessity suppose that C is not inferentially efficient for some $\tilde{y} = y$. Thus both the agent and the principal would prefer another control set C' to C. The principal prefers C' to C conditional on the agent's acceptance of C' over C. Since the agent could choose a decision in C (by rejecting C'), this implies that $C \cup C'$ is preferred by the principal to C. Since $C \subset C \cup C'$ the agent prefers $C \cup C'$ to C for all $\tilde{y} = y$. Thus C cannot be weakly efficient. This proves necessity.

Q.E.D.

Theorem 2.17 exhibits a certain stability property of weakly efficient control sets. Suppose the principal and the agent agree to use a certain control set C before the agent sees the outcome of $\tilde{y}$.

Then there will never be a reason for them to change to another control set after $\tilde{y}$ is revealed to the agent, since such a move cannot benefit both parties. On the other hand, it is clearly irrational for the principal and the agent to employ a control set which is not weakly efficient. In particular, it is immediate that the delegation solution is weakly efficient (provided it is unique).

If C' dominates C in terms of weak efficiency, it must be that $C \subset C'$ (provided C does not contain decisions, which the agent will not take under any outcomes of $\tilde{y}$, in which case C' could be augmented, if necessary, by those decisions without affecting the game). This fact makes it generally rather simple to identify weakly efficient control sets. For instance, in the production problem (example 2.2), the weakly efficient control intervals are of the form $[0, d_U + \varepsilon]$, where $\varepsilon > 0$ and d_U is the upper limit of the optimal control interval to the principal's problem.

Weak efficiency is very similar to the notion of conditional efficiency of outcome functions, proposed by Wilson [1977]. An outcome function d(y) is said to be <u>conditionally efficient</u> if there does not exist another outcome function d'(y), such that

$$E[F_i(d'(y),z)|y_i] \geq E[F_i(d(y),z)|y_i] \qquad \forall y_i, \forall_i,$$

with strict inequality for at least one i and some $y_i \varepsilon Y_i$.

The difference between weak efficiency and conditional efficiency is that in weak efficiency we require that d(y) should be attainable via the corresponding decision mechanism N = (d,Y), whereas conditional efficiency is purely a statement about an outcome function. Thus weakly efficient outcome functions may not be conditionally efficient. However, in Wilson's model it is assumed that $\tilde{y}$ is observable _ex post_, and thus any outcome function can easily be attained as a Nash equilibrium. In this sense our definition is an extension of Wilson's definition to cover situations in which $\tilde{y}$ is not observable _ex post_.

Following Wilson [1977], it can easily be shown that weakly efficient outcome functions (i.e., outcome functions corresponding to weakly efficient decision mechanisms) can be generated by solving:

$$\max_{d(\cdot)} E[\sum_{i=1}^{n} \lambda_i(y_i)\ E[F_i(d(y),z)|y_i],$$

$$\text{s.t.} \qquad d(y) \text{ attainable.}$$

The constraint that d(y) should be attainable, expresses the limitations that are imposed by the presence of differential information, and in this sense weak efficiency is a constrained efficiency notion.

2.6 Concluding Remarks

In Chapter I we formulated a general model of decentralized decision-making with the interpretation that agents participate by sending messages to the principal, who makes the final decision. In this chapter we have studied a special but important case of the general model, which we called delegation. Delegation was seen as a process where the agent is given freedom to act in his own self-interest within a constrained set of alternatives.

Delegation can take many forms. We chose to investigate quantity controls and the use of more general nonlinear price-schedules. The study of quantity controls allowed us to focus on the relationship between the key determinants of optimal delegation: differential information and differences in objectives. For the agent's preference structures a partial ordering was developed. The optimal amount of delegation as well as the value of the agent to the principal was a monotone function of this partial ordering. The results also indicated that an increase in the information gap was normally accompanied with greater freedom for the agent.

In the context of a centrally planned economy this can be interpreted as saying that the tightness of economic control depends on both the curvatures of the revenue and cost functions as well as on how informed the center is. If the center is badly informed less rigid controls should be used. This complements Weitzman's results on the impact of curvatures alone on economic control.

We also found that delegation, in addition to its role as an instrument of information transmission in decision-making, reflected motivation and learning aspects. We gave an explanation to the much held opinion that delegation motivates agents to perform better. From the point of view of economic planning, this gives prices a comparative advantage over quantities.

The applicability of delegation depends on the preferences of the agent as compared to the principal. There has to be a certain degree of conformity in their responses to the information signal. For quantity controls a simple sufficient condition was developed; if the agent takes acts both below and above the principal's best centralized choice, delegation leads to Pareto improvements. This condition may be quite generally satisfied. But if preferences do not conform as much as is required by this condition, prices can be used to change the agent's preferences so that delegation becomes valuable.

If interests lie even further apart, observations that are made ex post can be included in the delegation arrangement so as to allow successful delegation. An example of this is the goal-based incentive scheme (Example 2.4). Unless the outcome could be observed afterwards, there would be no point in letting the agent set his own goal. Tying the agent's reward to the outcome, makes his information (transmitted via the goal) credible. The same principle is at work, for instance, when warranties are used, particularly in markets for used commodities. Part of the rationale for extensive cost accounting can also be found in this principle, though we will discuss other reasons later (Chapter IV).

Changing preferences, making information credible, measuring output for extra signals; all are costly activities, which must imply (with few exceptions) that differential information causes inefficiencies by the standard of perfect information. We analyzed the issue of efficiency with differential information. When trying to apply the simple principle that Pareto optimality means that nobody can be made better off without hurting somebody else, we were led to introduce a noncooperative Nash equilibrium concept into the definition. We do not think this is very useful, but conclude from the analysis that it is plausible that there will be no generally viable definition of efficiency with differential information. Instead it may be that one has to find out how efficiency should be measured in the context of specific models or situations. Our own definition of weak efficiency should be seen in this light, and similarly Grossman's work on the efficiency of rational expectations equilibria [1977]. The same can be said about Wilson's notion of conditional efficiency, which seems best applicable in a situation where markets for signal outcomes are complete.

An important dimension of a principal-agent relationship is time. How long the agent is going to work for the principal is an essential question for determining the optimal control. This is particularly true when the principal has an imperfect model of the agent's response behavior. One observation we made, was that with an imperfect model and a multiperiod time span, the optimal control will reflect experimentation on behalf of the principal. When quantity controls are used, the principal will give the agent more freedom in order to learn the agent's characteristics faster.

This result assumed the agent will behave myopically or nonstrategically towards experimentation. Generally, it leads to a problem of adaptive control, which is by no means simple to handle (see Prescott [1972]). Even more difficulties are encountered if one assumes the agent behaves nonmyopically.

Another possibility is to move to a self-fulfilling expectations equilibrium directly. It can be envisioned as the agent telling the principal his characteristics at the outset. The principal selects his control optimally based on these characteristics, and then merely checks that the action-outcome pairs he can observe will be consistent with his beliefs about the agent. An important aspect is that the agent may choose to lie about his characteristics and then simulate a behavior consistent with this lie.

If we look at the production problem (Example 2.2) it is easily seen that the agent has an incentive to misrepresent his preferences so that they appear actually closer to the principal's. Both parties are better off with this arrangement, which reflects the fact that outcomes are used to validate the agent's information about his behavior. With transfer payments available, they could presumably reach a first-best solution (corresponding to perfect information).

Whether the analysis is so simple for misrepresentation of information is not clear. One would think, however, that the agent may have an incentive to appear less informed in order to cover self-interested actions, at least in some situations.

An issue which is familiar from other models of self-fulfilling equilibrium, is that there may be multiple equilibria, which can be Pareto ordered. This is a possibility when the agent's expertise is a function of his effort. We saw that the control size in this case determines the effort level, whereas the reverse is true when the principal looks for the control given the level of the agent's expertise. One could imagine that an agent could be inefficiently utilized, because the principal does not know the agent's potential.

APPENDIX 2A

EXISTENCE OF OPTIMAL CONTROLS

Recall our earlier assumptions:

A1. D is a compact subset of a complete, separable metric space.

A2. $\mathcal{C}$ is a closed subset of 2^D, w.r.t. the Hausdorff-metric (see below).

A3. F^A and F^P are continuous and uniformly bounded.

In proving existence, we will take the standard approach of showing that our problem is one of finding the supremum of a continuous function over a compact set. Since the argument of the objective function is a set, we have to define a suitable metric to get the objective function continuous. For this purpose we will use the Hausdorff (H-) metric (see Munkres [1975]), which is defined for two sets $A,B\varepsilon 2^D$ as

$$H(A,B) = \max \{\sup_{a\varepsilon A} \inf_{b\varepsilon B} m(a.b), \sup_{b\varepsilon B} \inf_{a\varepsilon A} m(a,b)\}$$

where m is the metric of the space containing D. The crucial result, which gives us compactness is the following (see Munkres [1975]):

Lemma 1: If D is a compact set in the metric m, then the set of all nonempty bounded closed subsets of D is compact in the Hausdorff-metric H.

Since we assumed $\mathcal{C}$ is closed in the H-metric, Lemma 1 tells us that $\mathcal{C}$ is compact, because 2^D is metrizable.

Lemma 2: $d(y,C)$ is u.s.c. in C for every y.

Proof: Fix y. Let

(i) $C_n \to C$ (in the H-metric), C_n, $C \varepsilon \mathcal{C}$

(ii) $d_n \varepsilon d(y,C_n)$

(iii) $d_n \to \hat{d}$.

We need to show that $\hat{d} \varepsilon d(y,C)$. $d(y,C)$ is nonempty, since C is compact. Let $\bar{d} \varepsilon d(y,C)$ be arbitrary. We show that $\hat{d} \varepsilon C$ and $f^A(\hat{d},y) = f^A(\bar{d},y)$, implying $\hat{d} \varepsilon d(y,C)$.

Let $\delta > 0$ be arbitrary. $C_n \to C \Rightarrow \exists n_1$ and a sequence $\{d(n)\}$, $d(n) \varepsilon C \ \forall n$, s.t. $m(d(n),d_n) < \delta$ when $n \geq n_1$. $d_n \to d \Rightarrow \exists n_2$ s.t. $m(\hat{d},d_n) < \delta$ when $n \geq n_2$. From the triangle inequality we get $m(d(n),\hat{d}) < 2\delta$ when $n \geq n_0 = \max(n_1,n_2)$. Since δ was arbitrary, $\hat{d}$ is a limit point of a sequence in the closed set C and so $\hat{d} \varepsilon C$.

$\hat{d} \varepsilon C \Rightarrow f^A(\hat{d},y) \leq f^A(\bar{d},y)$. If moreover $f^A(\hat{d},y) \geq f^A(\bar{d},y)$ we are done. $C_n \to C \Rightarrow \exists$ a sequence $\{d'(n)\}$, $d'(n) \varepsilon C_n$, $\forall n$, s.t. $d'(n) \to d$. We have that $f^A(d'(n),y) \to f^A(\bar{d},y)$, since f^A is continuous. We also have $f^A(d_n,y) \to f^A(\hat{d},y)$ by (iii). $d'(n) \varepsilon C_n$ implies $f^A(d'(n),y) \leq f^A(d_n,y) \ \forall n$, by (ii). Taking limits on both sides gives $f^A(\bar{d},y) \leq f^A(\hat{d},y)$ concluding our proof. Q.E.D.

Lemma 3: If Prob{d(y,C) is not a singleton} = 0 $\forall C\varepsilon\mathcal{C}$, then $E_y\{f^P(d(y|C),y)\}$ is continuous in C for any specific choice of response function from the response correspondence. In particular it is true for the choice $d_{min}(y|C)$.

Proof: Let $C_n \to C$. Let $\Lambda = \{y|d(y|C)$ is not a singleton$\}$. By assumption $P(y\varepsilon\Lambda) = 0$. $E_y\{f^P(d(y|C_n),y)\} = E_y\{f^P(d(y|C_n),y);\ y\varepsilon\Lambda\} + E_y\{f^P(d(y|C_n),y);\ y\varepsilon\Lambda^c\}$.

The first term in the RHS is 0. Let us show that $f^P(d(y|C_n),y) \to f^P(d(y|C),y)$ for $y\varepsilon\Lambda^c$. Then, since f^P is continuous and bounded, the integral will converge to the desired limit by the bounded convergence theorem.

Write $d_n = d(y|C_n)$; $d_n \varepsilon d(y,C_n)$. We claim $d_n \to d(y|C)$. This is true if and only if every subsequence $\{d_{n'}\}$ has a refinement $d_{n''} \to d(y|C)$. Since D is compact, any subsequence $\{d_{n'}\}$ has a convergent subsequence $d_{n''} \to \hat{d}$. Since d(y,C) is u.s.c. by Lemma 2, $\hat{d}\varepsilon d(y,C)$, which is a singleton (because $y\varepsilon\Lambda^c$) and so $\hat{d} = d(y|C)$. Hence, $d_n \to d(y|C)$. This completes the proof. Q.E.D.

Theorem 3.1: Assume A1-A3 and in addition Prob{d(y,C) is not a singleton} = 0 $\forall\ C\varepsilon\mathcal{C}$, then there exists an optimal solution $C^*\varepsilon\mathcal{C}$ to the delegation problem regardless of the maximizing response function the agent uses.

Proof: The delegation problem is to find a maximizing $C^* \varepsilon \mathcal{C}$, if it exists, for the function $E_y\{f^P(d(y|C),y)$, where $d(y|C)$ is a particular maximizing response function of the agent. By Lemma 1 $\mathcal{C}$ is compact and by Lemma 3 the objective function is continuous in C. This implies there exists an optimal control $C^* \varepsilon \mathcal{C}$. Q.E.D.

Let us now study the particular response function $d_{max}(y|C)$.

Lemma 4: If (i) $C_n \to C$

(ii) $E_y\{f^P(d_{max}(y|C_n),y)\} \to \bar{E}$

then $\bar{E} \leq E_y\{f^P(d_{max}(y|C),y)\}$. In other words, $E_y\{f^P(d_{max}(y|C),y)\}$ is u.s.c. in C.

Proof: Write $d_n = d_{max}(y|C_n)$, $E_n = E_y\{f^P(d_n,y)\}$

(A.1)
$$\bar{E} = \lim_{n\to\infty} E_n = \lim\sup E_y\{f^P(d_n,y)\} \leq \lim_{k\to\infty} E_y\{\sup_{n\geq k} f^P(d_k,y)\}$$
$$= E_y\{\lim\sup f^P(d_n,y)\}$$

(the inequality holds for sup's and taking limits, which exist since the sequences are decreasing, we get the limiting inequality; the last equality follows by bounded convergence). We claim: $\lim\sup f^P(d_n,y) \leq f^P(d_{max}(y|C),y)$ for each y.

It is easy to construct a subsequence $\{d_{n'}\}$ s.t. $f^P(d_{n'},y) \to \limsup f^P(d_n,y)$. Since D is compact there is a converging refinement $d_{n''} \to \hat{d}$. Of course, $f^P(d_{n''},y) \to \limsup f^P(d_n,y)$. Since f^P is continuous, we also have $f^P(d_{n''},y) \to f^P(\hat{d},y)$. Hence $\limsup f^P(d_n,y) = f^P(\hat{d},y)$. By Lemma 2 $d(y,\cdot)$ is u.s.c. for every y. This implies $\hat{d}\varepsilon d(y,C)$ and so $f^P(\hat{d},y) \leq f^P(d_{max}(y|C),y)$ by definition of the d_{max}-function. This proves the claim and the lemma follows directly from (A.1).

Q.E.D.

Theorem 2: For the response function $d_{max}(y|C)$, there exists an optimal control $C^*\varepsilon\mathcal{C}$.

Proof: The theorem is a direct consequence of the u.s.c. of $E_y\{f^P(d_{max}(y|C),y)\}$ and the compactness of $\mathcal{C}$ (see Luenberger [1968]).

Q.E.D.

APPENDIX 2B

THE HAZARD RATE FOR A NORMAL DISTRIBUTION

The hazard rate is defined as:

$$h(y) = \frac{g(y)}{1 - G(y)}, \quad y \sim N(0,s^2).$$

(B.1) $$h'(y) = \frac{g'(y)}{1 - G(y)} + \frac{g^2(y)}{(1 - G(y))^2} = h(y)(h(y) - y).$$

(B.2) $$h''(y) = h(y)[(h(y)-y)^2 + h'(y) - 1].$$

It is well known that $h'(y) > 0$ (see Barlow and Proschan [1975]). Consequently, (2.16) follows from (B.1).

Let ψ,Φ be the density and distribution functions of a standardized Normal distribution. Using ℓ' Hospital's rule twice we find that:

(B.3) $$\frac{\psi(y)}{y \cdot (1 - \Phi(y))} \to 1, \text{ as } y \to \infty. \quad \text{Hence,}$$

$$\frac{g(y)}{y(1 - G(y))} = \frac{\psi(y/s) \cdot \frac{1}{s}}{y(1 - \Phi(y/s))} \to \frac{1}{s^2}, \text{ as } y \to \infty,$$

which is (2.18).

Let h_0 be the hazard rate for a standardized Normal distribution. Suppose $h_0'(\bar{y}) > 1$ for some $\bar{y}$, contrary to the claim in (2.17).

By (B.2) $h_0''(\bar{y}) > 0$, implying $h_0'(y) > 1$ for all $y > \bar{y}$. This contradicts (B.3), since $h_0(y) > y$ for all y. Hence, $h_0'(y) < 1$ for all y. (2.17) follows then from the fact that $h'(y) = h_0'(y); 1/s^2$, since we already stated that $h'(y) > 0$.

To prove (2.19) we have:

$$\frac{\partial}{\partial s}[h(y) \cdot s^2] = \frac{\partial}{\partial s}[h_0(y/s) \cdot s] = h_0(y/s) + s \cdot h_0'(y/s) \cdot (-y/s^2).$$

If $y \leq 0$, this expression is certainly > 0, and when $y > 0$, we can minorize it by using (2.16) and (2.17) and conclude:

$$\frac{\partial}{\partial s}[h(y) \cdot s^2] > y/s - s \cdot 1 \cdot \frac{-y}{s^2} = 0.$$

This establishes (2.19).

Footnotes to Chapter II

[1]A word about terminology. Asymmetric information is, of course, present in classical models of price-mediated markets. What are unknown are the preferences and endowments of individuals. When we talk about asymmetric (or differential) information we have in mind information which is of interest to other agents in the economy. This is not assumed to be the case in classical models, since preferences are independent.

[2]For instance, in Spence's model of labor markets, individuals purchase education in excess of what would be efficient under symmetric information.

[3]For notational convenience we will include the agent's productive inputs in D, since these do not play a central role in the analysis.

[4]$\tilde{y}$ may pertain both to information about the underlying problem that the principal wants to solve, and to characteristics of the agent which the principal does not know with certainty; e.g., the agent's level of expertise and his preferences. If the principal knows the agent's characteristics fully, we say that the model is perfect. Otherwise it is imperfect. Though this distinction is not stressed in the sequel, the assumptions in Section 2.3 make sense only in a perfect model. However, they can be easily modified to yield the same conclusions in imperfect models (with few exceptions), and this will be pointed out later.

[5]We will only be concerned with perfect Nash equilibria (see Selten [1974]), and omit considerations of possible threats from the agent.

[6]This theorem guarantees the existence of a perfect Nash equilibrium. See footnote 5.

[7]The range of $\tilde{z}$ is assumed to be such that the cost is an increasing function in d, except when we use the Normal distribution as an approximation (see later analysis of the example in Section 2.3.2).

[8]The goal-based incentive scheme described above is an example of a scheme which simultaneously serves the purpose of providing incentives for productive inputs and truthful communication of information. For instance, the goal may be taken as a signal for production potential when making investment decisions (cf. Weitzman [1976a]).

[9]For instance, if the agent is a firm whose pollution level is to be controlled, the installment of a filtering device may guarantee that the level of pollution lies in an acceptable range. No measurement of the pollution level is necessary, as would be the case if a price scheme was used.

[10]$d(r) = \infty$ is a distinct possibility unless we assume U is unbounded. It will only strengthen the point we are making, namely that quantity controls are needed.

[11]The Normal distribution is necessarily an approximation when we use it in the production problem, since we want the cost function $C(d,z) = q^2 - 2 \cdot z \cdot d$ to be increasing in d for all z. Similarly, some responses in (2.13) will be negative, which is inconsistent with the interpretation that d is a production decision. One could work with a truncated Normal distribution instead, but this would lead to unnecessarily complicated algebra.

[12]By the best centralized act, we mean the decision that the principal would take if he would not have the agent available.

[13]Under appropriate convexity assumptions, of course.

[14]These approximations can be defended rigorously using the results in Samuelson [1970]. The assumption is that s is sufficiently small.

[15]To mention a few: Ireland [1977], Laffont [1977], Spence [1977], Yohe [1977a], [1977b].

[16]Of course, the analysis of controlling the demander is symmetric to the one given above, but Laffont addresses the issue of which side should be controlled by prices, which side by quantities (or maybe both by quantities).

[17]Compare this to the recent literature on fixed price equilibria (see for instance Benassy [1975]).

[18]Of course, when s gets large, the quadratic approximations loose their validity and the results have to be viewed in terms of an example only.

[19]As we pointed out in footnote 4, this theorem is true for both imperfect and perfect models, but the coherence condition as stated

is only meaningful in a perfect model. An extended coherence definition would state that each possible response function of the agent is coherent. Theorem 2.6 would still be true under this weaker coherence condition as is easily seen from the argument in the proof.

[20]Let $f_1(y)$ and $f_2(y)$ be two density functions such that f_2 is a mean-preserving spread of f_1. Let $\bar{y}$ be the mean for both distributions. Then, by definition of a mean-preserving spread,
$\int_y^\infty f_1(y)dy < \int_y^\infty f_2(y)dy$, for every $y > \bar{y}$. This implies
$\int_y^\infty g(y)f_1(y)dy < \int_y^\infty g(y)f_2(y)dy$, for every $y > \bar{y}$, whenever $g(y)$ is an increasing function.

[21]Let us stress that this result refers to a perfect model or to information about the principal's problem. If the principal gets more uncertain about the agent's characteristics, the implications are quite the opposite. The agent will be given less freedom when the principal knows less about him. Compare this with the discussion in Section 2.3.8.

[22]One may ask if Theorem 2.9 would have been valid if uniform closeness had been defined as follows: A is uniformly closer in preferences than A' if, for each y, P prefers $d_A(y)$ to $d_{A'}(y)$. Clearly, the assertion that A' is given a control which is included in A:s control set, is not generally true. For instance, if $b_A > b_P > b_{A'}$, in the production problem (Example 2.2), the optimal control for A has the form $(-\infty,\bar{d})$, but for A', $(\underline{d},\infty)$. More surprisingly maybe, it is not even true generally that the principal is better off with an agent whose uncontrolled response function he prefers pointwise. Unless the agents' response functions lie on the same side of the principal's response function (as we have required), controls may change the pointwise closeness in favor of the agent whose uncontrolled response was worse.

[23]Again, the coherence condition could be weakened as discussed in footnote 19.

[24]If effort improves the outcome directly, this does not explain why prices would be preferred to quantities. For the agent's self-interest he would provide equal amounts of effort to improve production regardless of whether he made the decision on the level of production or if it was the principal who did it.

[25]This, of course, includes the processing of information as well as collecting it. It is not what he has in his portfolio, but what he has in his head that we count as information.

[26]Note, though, that this result is not necessarily valid if the extra effort on information gathering will reduce the manager's effort expenditure on supervision and implementation of the decision.

[27]A Nash-equilibrium of the one-period problem remains a Nash-equilibrium in the multiperiod problem.

[28]See Prescott [1972] or Grossman, Kihlstrom and Mirman [1977].

[29]As we saw in the section on interval controls, intervals are generally not optimal. This has the implication that the optimal solution to the regulation problem may quite easily be nondifferentiable, which further complicates the analysis.

[30]This part of the proof is found in Riley [1976].

[31]For a more detailed discussion, see Marschak and Radner [1972].

[32]It is clear that the traditional definition is a special case of inferential efficiency.

[33]Recall from Chapter I that there is no loss in generality to take the message space equal to Y, the range of the agents' signals.