# Delegation and discretion

Armstrong, Mark

University College London (UCL)

1995

# Delegation and Discretion

Mark Armstrong
Department of Economics
University of Southampton

November 1995

**Abstract**

There are many situations in which a principal delegates decisions to a better-informed agent but does not choose to give full discretion. This paper discusses one reason why this might be desirable: the agent may have tastes that differ from those of the principal. Limiting the agent's discretion has the advantage that an untrustworthy agent is constrained from following policies that are disliked by the principal, but the disadvantage that trustworthy agents are then not permitted to carry out some desirable policies. It is shown that a greater risk of the agent being untrustworthy will lead to her being offered less discretion over policy. Applications of the model involve judicial sentencing policy, monetary policy, and pricing policy in a regulated industry.

**Keywords:** Principal-Agent Problem; Delegation; Discretion; Mandatory Sentences; Monetary Policy; Regulation.

# 1 Introduction

There are many situations in which a principal delegates decisions to a better-informed agent, but does not choose to give the agent full discretion.[1] Examples discussed in this paper include the imposition of mandatory sentences on the judiciary, the limits placed on central bankers concerning monetary or exchange rate policy, and the limits placed on industry regulators concerning pricing policy. This paper proposes a framework in which to analyse these and similar situations.

There are a variety of reasons why it may be desirable to limit agency discretion, and they usually involve the possibility that the agent may not always be trusted to implement the principal's ideal policy. This presents an obvious danger for delegation, and the principal must trade off the benefits of offering wide agency discretion — the ability of trustworthy agents to be able to respond freely to changing circumstances — against the drawbacks — the ability of an untrustworthy agent to pursue undesirable policies.

A trivial example may help to fix ideas. Suppose a government is considering policy towards the speed of vehicles on roads. The ideal maximum speed of a given driver on a given road is a function of, *inter alia*, the traffic conditions on the road, the weather, the type of road, the type of car, and the skill of the driver. While it is commonplace to link the speed limit to the type of road, it is impractical to link the permitted speed to any other of the above variables. On the other hand, a given driver *can* condition her speed on each of these factors, but she may have different views on the desirable speed given these factors compared to the government. The government then faces an obvious tradeoff: granting drivers full discretion allows for the full use of private information in determining speed, but also allows drivers with different preferences (e.g. reckless drivers) to drive at speeds which the government feels inappropriate given the conditions. However, although imposing a speed limit does mitigate the danger of, say, reckless driving, it also prevents perfectly safe drivers from driving fast should conditions be appropriate. Presumably, the greater the proportion of drivers whose preferences over travel time and safety coincide with the government's, the more discretion drivers as a whole should be given — see Proposition 1 below.

[1] An early version of this paper was presented at a joint CEPR/ECARE workshop on 'The Political Economy of Competition Policy' in Brussels, November 1993. I am grateful to John Driffill, Bruno Jullien, Paul Klemperer, Tracy Lewis, Preston McAfee, Ray Rees, Patrick Rey, David Salant and Paul Seabright for helpful comments.

In general, there are at least three reasons why an agent might not pursue correct policies should they be given discretion:

1. The first is the problem of dynamic consistency which has been much analysed in a variety of contexts.[2] In the regulation setting a regulator may be tempted to set low prices — prices that cover only the firm's future average avoidable costs — once sunk-cost investments on the part of the firm have been made. One way to overcome this danger of *ex post* opportunism might be for government to limit the regulator's discretion by imposing the rule that the firm always be granted a fair return on 'used and useful' capital.[3]

2. Then there is the possibility of the agent being influenced by various interest group pressures. This 'political economy' or 'capture' theory of decision-making was developed by Stigler (1971), Posner (1974) and Peltzman (1976), and these papers argue that limits to agency discretion are desirable, both in order to prevent bad policies being followed by captured agents, and to remove the incentive for interest groups to engage in wasteful rent-seeking behavior.

3. As discussed in the speed limit example above, a third reason is that the agent may have fixed but divergent preferences over policy from those of the principal.[4] This could be seen as the opposite polar case to that of capture. Instead of modelling the problem as one in which all potential agents have identical and known utility functions but who may be tempted to alter their actions by inducements from interested parties, here agents' tastes are exogenous but unknown. Take the case of regulation. Conditions within the industry concerned are uncertain — for instance, the cost level or the potential for cost reduction are things that are not known *ex ante* — and the regulator presumably is better informed about these conditions than government. (That is one reason why governments employ specialist regulators.) For this reason the government will delegate decisions about, say, the tightness of price control to the regulator. This system works perfectly well provided government and regulator agree on the correct policy given circumstances within the industry; that is to say, if their preferences coincide. However, it is not always reasonable to assume that this is so, and the regu-

[2]See, among many others, Kydland and Prescott (1977) and Gilbert and Newbery (1994).

[3]For a discussion of this point, see Greenwald (1984).

[4]Vickers (1985) examines a model where a principal *deliberately* hires an agent with divergent tastes in order to be able credibly to commit to desirable policies. For instance, in a Cournot oligopoly setting, if a firm's owners (who are interested in profit) hire a manager who is also interested in the firm's output (and whose preferences are publically known), then in equilibrium the firm's profits could increase. For a similar idea in the monetary policy area, see Rogoff (1985).

lator may tend to favour consumer or industry interests to a greater extent than the government would wish and for this reason it may be desirable to limit regulatory discretion. This reason for limiting discretion has received less attention than the others, and it is the purpose of this paper to discuss some of the issues involved.

One recent paper which does not fit into any of these categories is the study of the allocation of decision-making within organizations by Athey *et al.* (1994). In their model there is a continuum of possible types of decision which might need to be made, indexed by a random variable $s \in [0, 1]$, and there are two people who can take any decision, a senior or a junior agent. (One application is to doctors and nurses, where the question is which profession should treat which illnesses.) The authors assume that the senior agent has a comparative advantage in making the decisions labelled with high $s$ (the important decisions). As a result, the junior person makes all decisions in a range $[0, s^*]$ and the senior person makes the remaining decisions. The primary purpose of the paper is to analyze what determines $s^*$, the degree of discretion of the junior agent. There are no information asymmetries in the model, and the senior person is well-informed about both the state $s$ and the ability of the junior person. The reason that the allocation of decision-making is not simply made on the basis of which agent has the absolute advantage in making the particular decision is because it is assumed that there are decreasing returns to decision making, and the more decisions the senior agent has to make the lower the quality of her decisions. Because of this, at the marginal decision (the highest state in which the junior agent makes the decision, denoted $s^*$ above) it may be that the senior agent has the absolute advantage. (The authors interpret this as a kind of 'hands-off' management style.) They also show that as the likelihood of more the important decisions increases, the junior agent should be given more discretion over policy (in the sense that $s^*$ increases), although the effect on the *frequency* of junior decision-making is ambiguous.

The most closely related paper to this one, however, is Holmstrom (1984). From section 3 onwards in his paper he discusses a model which is very similar to that discussed here in sections 2 and 3. In particular, he has a model where a principal delegates a decision to a better-informed agent who may have different preferences over policy. As in the present paper, the principal controls the agent by restricting the discretion over policy granted to the agent, rather than by using monetary incentives for instance. In section 5 he assumes that discretion is restricted by offering the agent an *interval* of choices from which she is free to choose. In addition he has a result (Theorem 3 in his paper) that is similar to the main result of this paper, namely, that an agent with preferences close to the principal's will be given more discretion than one with more divergent preferences. The main difference between the two papers is that the agent in Holmstrom's framework has known preferences, whereas in this paper the principal is uncertain about the agent's preferences. Another difference is that I use elementary mathematics throughout, and try to emphasize the various applications of this sort of model.[5]

The plan of this paper is as follows. A simple model is presented in section 2 in which the principal is uncertain of the decision the agent would like to make. There it is shown that the more *reliable* the agent is, the more discretion over policy she should be given. The modelling of the reliability of the agent is given one possible foundation in section 3, where the agent is assumed to have divergent tastes from the principal. Applications of this theory are discussed in section 4, and these are to sentencing policy, regulatory pricing policy, and to monetary policy. Finally, conclusions and possible directions for further work are discussed in section 5.

# 2 A Model

A *Principal* is planning to delegate the making of a decision to an *Agent* (or a population of agents). The Principal gains utility $u(d, \theta)$ if decision $d$ is made when $\theta$ is the state of the world. Both $d$ and $\theta$ are assumed to be scalars.[6] The Principal aims to maximize the expected value of this utility function. The decision variable $d$ is constrained to lie in the real interval $[d_{\min}, d_{\max}]$. Utility $u(\cdot, \theta)$ is concave in $d$ and the marginal utility of increasing $d$, $u_d(d, \theta)$, is increasing in $\theta$. The Principal is assumed to be unable to observe $\theta$ (otherwise she could simply maximize her utility given $\theta$ and obtain the first best), but believes $\theta$ to be distributed with support $[\theta_{\min}, \theta_{\max}]$ and density function $f(\theta)$. Let $d(\theta)$ be the first-best choice for

[5] Another related paper is De Bijl (1995) where an agent has to choose between a finite number of projects over which she has unknown preferences from the point of view of the principal. The agent can increase the principal's payoff by expending effort, the cost of which is negatively related to the agent's private benefits of the project (which is private information). As in this paper and Homstrom's paper, De Bijl assumes that the principal does not attempt to influence the agent's decision using monetary incentives, but rather by simply restricting her options.

[6] Although it is crucial in all the following analysis that the decision variable $d$ be a scalar, it is relatively straightforward to allow $\theta$ to be multi-dimensional. Indeed, in Armstrong (1994) and Holmstrom (1984) $\theta$ is allowed to be a vector. However, it does simplify the analysis somewhat — notably in Lemma 1 below — and so in this paper $\theta$ is a scalar.

the Principal given $\theta$, i.e.

$$d(\theta) \text{ maximizes}_{d_{\min} \leq d \leq d_{\max}} \ u(d, \theta) \ .$$

Because $u_d$ is increasing in $\theta$, it follows that $d(\theta)$ is increasing in $\theta$. Finally, let $d^*$ denote the Principal's best choice in the absence of delegation, i.e.

$$d^* \text{ maximizes}_{d_{\min} \leq d \leq d_{\max}} U(d)$$

where $U(d) = E_\theta u(d, \theta)$ is the Principal's expected welfare if the policy $d$ is followed in all states of the world $\theta$, and $E_\theta[\cdot]$ takes expectations over $\theta$. (Holmstrom (1984) calls $d^*$ the optimal *centralized* decision.) Since $u$ is concave in $d$, if follows that $U$ is also concave in $d$.

The Agent, unlike the Principal, is assumed to observe $\theta$. The Agent, if unconstrained, is assumed to follow a policy given $\theta$ which the Principal believes to be stochastic. In particular, if the state of the world is $\theta$, the probability that the unconstrained Agent chooses a policy $\delta \leq d$ is given by the distribution function $G(d, \theta)$. This function summarises all the relevant information about the Agent. For now, we will analyze the Principal's optimal policy using this 'reduced-form' behaviour function for the Agent; later in section 3 we will derive the function $G$ using more primitive assumptions about the Agent's preferences. For simplicity, suppose that $G$ is a smooth function, and write $g(d, \theta) = G_d(d, \theta)$ to be the density function for $d$ given $\theta$. Assume that $G(d_{\min}, \theta) = 0$ and $G(d_{\max}, \theta) = 1$ for all $\theta$. Most importantly, I assume that, like the Principal, the Agent is more likely to choose a higher value of $d$ if $\theta$ is higher; in other words

$$G(d, \theta) \text{ is decreasing in } \theta. \tag{1}$$

This assumption — that the Principal and Agent agree about which values of the parameter $\theta$ should lead to higher values of $d$ being chosen — plays the role of Holmstrom's assumption that preferences are *coherent* (Holmstrom, 1984, p.127). Without this assumption it is likely that delegation will have no benefits and the Principal will be able to do no better than to follow policy $d^*$ in all states of the world — see footnote 9 below.

Because of the danger of the Agent's ideal choice diverging from that of the Principal, the latter wishes to influence the decision of the former. In theory, one method of doing this is for the Principal to make the income of the Agent depend in some way on her decision. For instance, if the Agent is likely to choose a lower value of $d$ than the Principal would wish — so that $G(d(\theta), \theta)$ is rather large — then the Principal may well wish to make the pay of the Agent depend positively on the value of $d$ chosen in order to counteract this tendency. However, I make the strong assumption that the Principal chooses *not* to use the instrument of decision-contingent payments (or is constrained not to do so). Instead, she influences the Agent's behavior only by means of restricting the set of decisions the Agent may choose. Thus, rather than being permitted to choose $d$ from the whole set of possible decisions, $[d_{\min}, d_{\max}]$, the Agent is constrained to choose from some strict subset of decisions, say $D \subset [d_{\min}, d_{\max}]$.

There are several reasons why I make this assumption. First, in practice it is often the case that when delegating decisions to agents, principals do *not* make choice-contingent payments. This is certainly true for the applications discussed below in section 4. In this sense, the analysis of the paper is rather descriptive and addresses the sub-problem of how much discretion to give the agent *given* that there is no incentive scheme in place (or more generally, given a fixed incentive scheme), rather than aimed at the more fundamental question of why such incentive schemes are *not* observed in practice in many cases.

A second possible justification for assuming away decision-related payments could be the one discussed in Holmstrom and Milgrom (1991). They analyze a setting in which an agent has several tasks to perform, each of which competes for her time, and her performance in some of these tasks is not observable to the principal. Therefore, by making her income depend positively on those dimensions of performance that *are* observable, there is the danger that the agent will concentrate her time on the more lucrative tasks to the detriment of the others, which, although unobserved, are beneficial for the principal. As a result, it may be optimal not to use performance-related incentive payments in a multitask setting. Although their model involves moral hazard rather than adverse selection, similar arguments could be made within the framework of this paper. Looking ahead to section 4.1 below, for instance, if it were generally felt that judges were on average too lenient in their sentencing policy for a particular crime then one option might be to reward judges on the basis of their sentencing record. However, together with other drawbacks, such a scheme could distort a judge's incentives over other, unobserved decision variables (such as whether to rule on debatable points of law in favour of, or against, the defendant).[7,8]

[7] Another example of multitask distortions in the legal context is the following: a proposal in Britain which surfaces from time to time is to make the likelihood of a judge being promoted to a more senior position depend negatively upon the number of their cases in which they find against the defendant and which are over-turned on appeal (and without any new evidence being produced). If such cases are penalized severely, this scheme could give an incentive for judges to find *for* the defendant more often than otherwise.

[8] Another possible theoretical justification for not using decision-contingent payment

A third way to justify the exclusion of decision-contingent payment schemes is simply to assume that the agent is 'infinitely risk-averse' and so cares only about her minimum possible payment. In such a case there is nothing to be gained by the principal in offering monetary incentive schemes. This modelling approach is used in Aghion and Tirole (1994), for instance.

Finally, the analysis required to determine the optimal payment schedule would in general be rather demanding. For instance, in section 3 below, the model involves two dimensions of informational asymmetry (the taste parameter $\alpha$ and the state of the world $\theta$), and this implies we would have to enter the difficult area of multidimensional mechanism design.

Suppose, then, that the Principal influences the behavior of the Agent by means only of limiting her discretion to some subset $D$ of all possible choices. I make the further simplifying assumption that the Principal considers only sets $D$ that are *intervals*:

**Assumption**: *The Principal considers only discretion sets of the form*

$$D = [d_1, d_2] \ \subset [d_{\min}, d_{\max}] \ . \tag{2}$$

I do this partly because it seems a natural way to represent the degree of 'discretion' given to the Agent and, again, partly because the general case of allowing *any* (compact) set $D \subset [d_{\min}, d_{\max}]$ is technically difficult.[9] Thus the Agent is free to make her choice of policy provided her choice $d$ lies between the lower bound $d_1$ and the upper bound $d_2$. The Agent is given no discretion if $d_1 = d_2$ (in which case it is optimal for the Principal to set $d_1 = d_2 = d^*$).

The function $G(d, \theta)$ describes the Agent's behaviour given no constraints. We have to make an assumption about her behaviour given bounds on her discretion, and I make the simplest assumption, namely, if $d$ is the Agent's optimal choice (given some state of the world $\theta$) and this lies outside the permitted range, then she will choose the point in the permitted range which is closest to her ideal choice, i.e. if $d < d_1$ then she will choose $d_1$ and if $d > d_2$ she will choose $d_2$. (Implicit in this is an assumption that the Agent's preferences are *single-peaked* in $d$ for each $\theta$.)

Given this assumption it is straightforward to calculate the Principal's expected utility given the bounds $d_1$ and $d_2$. First of all, fix the state of the

schemes is given by Laffont and Tirole (1993, chapter 11), who show it is possible that choice-contingent payments to the agent are undesirable *ex ante*. That is to say, the abilty of the principal to make such payments to the agent could leave the *principal* vulnerable to capture by interested parties (including the agent), and in some case it is optimal to constrain the principal from using such incentive payments.

[9] Holmstrom (1984, Appendix A) proves under fairly general conditions that there exists an optimal compact set of options for the Principal, but he has no results concerning when this optimal set is an interval. As in this paper, Holmstrom (1984, section 4 onwards) makes the *ad hoc* assumption that the control sets are intervals.

world $\theta$. Then the Principal's expected utility given $\theta$ and the bounds $d_1$ and $d_2$, denoted $w(d_1, d_2, \theta)$, is given by

$$w(d_1, d_2, \theta) = G(d_1, \theta)u(d_1, \theta) + [1 - G(d_2, \theta)]u(d_2, \theta) + \int_{d_1}^{d_2} u(\delta, \theta)g(\delta, \theta)\, d\delta\ . \tag{3}$$

(The first term gives the Principal's utility when the lower bound $d_1$ is binding for the Agent, which occurs with probability $G(d_1, \theta)$, the second is her utility when the upper bound binds, while the third is her expected utility when the Agent is not constrained by the bounds.) This can be more conveniently be written in the separable form

$$w(d_1, d_2, \theta) = \phi_1(d_1, \theta) + \phi_2(d_2, \theta) - \phi_0(\theta) \tag{4}$$

where

$$\begin{aligned}
\phi_0(\theta) &= \int_{d_{\min}}^{d_{\max}} u(\delta, \theta)g(\delta, \theta)\, d\delta \\
\phi_1(d, \theta) &= G(d, \theta)u(d, \theta) + \int_{d}^{d_{\max}} u(\delta, \theta)g(\delta, \theta)\, d\delta \qquad (5) \\
\phi_2(d, \theta) &= [1 - G(d, \theta)]u(d, \theta) + \int_{d_{\min}}^{d} u(\delta, \theta)g(\delta, \theta)\, d\delta\ .
\end{aligned}$$

In words, the function $\phi_0$ gives the Principal's expected welfare given $\theta$ if the Agent is given *full* discretion over policy, $\phi_1$ is expected welfare if she sets only a *lower* bound on the Agent's discretion given by $d$, while $\phi_2$ gives her expected welfare if she sets only an *upper* bound given by $d$. Finally, the Principal's expected welfare over all states of the world $\theta$, denoted $W(d_1, d_2)$, is just the expectation of $w$:

$$W(d_1, d_2) = \Phi_1(d_1) + \Phi_2(d_2) - \Phi_0 \tag{6}$$

where $\Phi_0 = E_\theta \phi_0(\theta)$ is the Principal's expected utility over all states of the world obtained by granting the Agent full discretion over policy, $\Phi_1(d) = E_\theta \phi_1(d, \theta)$ is her expected utility from setting only a lower bound of $d$, and $\Phi_2(d) = E_\theta \phi_2(d, \theta)$ is her expected utility from setting only an upper bound of $d$. The optimal bounds on the Agent's discretion, then, are given by maximizing $W$ with respect to $d_1$ and $d_2$ subject to the constraint $d_1 \leq d_2$. From (5), since $\partial\phi_1(d, \theta)/\partial d = G(d, \theta)u_d(d, \theta)$ and $\partial\phi_2(d, \theta)/\partial d = [1 - G(d, \theta)]u_d(d, \theta)$, we obtain

$$\begin{aligned}
\Phi_1'(d) &= E_\theta \left[G(d, \theta)u_d(d, \theta)\right] \qquad (7) \\
\Phi_2'(d) &= E_\theta \left[[1 - G(d, \theta)]u_d(d, \theta)\right]
\end{aligned}$$

These expressions give the marginal benefit of increasing the lower and upper bounds, respectively.

At first glance it might be thought that we would have to impose the constraint that $d_1 \leq d_2$ when maximizing $W$ in (6). However, the next result shows this not to be so.

**Lemma 1**

*(i)* $\Phi_1(d)$ *is decreasing for* $d > d^*$;

*(ii)* $\Phi_2(d)$ *is increasing for* $d < d^*$.

**Proof.** See Appendix.

This has the corollary that if $d_1$ maximizes $\Phi_1(d)$ then $d_1 \leq d^*$, and if $d_2$ maximizes $\Phi_2(d)$ then $d_2 \geq d^*$. In particular, the optimal lower bound on the Agent's discretion maximizes $\Phi_1(d)$ while the optimal upper bound maximizes $\Phi_2(d)$, and we need not worry about the constraint $d_1 \leq d_2$. Finally, since $d^*$ is the optimal policy without delegation and $d_1 \leq d^* \leq d_2$, the optimal interval offered to the Agent necessarily includes this optimal non-delegation decision. Here, if the Principal sets $d_1 = d_2 = d$ the Agent is given no discretion, and so $U(d) \equiv \Phi_1(d) + \Phi_2(d) - \Phi_0$. This is maximized at $d = d^*$. The two functions $\Phi_i(\cdot)$ are maximized at $d = d_i$ which necessarily bracket $d^*$. Moreover, we must have $\Phi_1(d_{\min}) = \Phi_0$ since if only a lower bound is set and set equal to the *minimum* possible $d$, then the Agent is effectively given full discretion. Similarly, $\Phi_2(d_{\max}) = \Phi_0$. Finally, it follows that $\Phi_1(d_{\max}) = U(d_{\max})$ since if a only a lower bound is set and set equal to the *maximum* possible $d$, then the Agent is effectively granted no discretion and expected utility is just $U(d_{\max})$. Similarly, $\Phi_2(d_{\min}) = U(d_{\min})$.[10]

The next Lemma describes the case of 'one-sided' errors:

**Lemma 2**

*(i) If* $G(d(\theta), \theta) = 0$ *for all* $\theta$*, then* $\Phi_1(d)$ *is decreasing for all* $d$;

*(ii) If* $G(d(\theta), \theta) = 1$ *for all* $\theta$*, then* $\Phi_2(d)$ *is increasing for all* $d$.

**Proof.** Obvious from proof of Lemma 1.

[10] It seems worth at this point to stress how important assumption (1) is if delegation and discretion is to be of any value to the Principal. Suppose, in contrast to (1), that $G(d, \theta)$ was *increasing* in $\theta$, so that the Agent was more likely to choose a low value of $d$ if $\theta$ is high. In this case the proof of Lemma 1 shows that (i) $\Phi_1(d)$ is increasing for $d < d^*$, and (ii) $\Phi_2(d)$ is decreasing for $d > d^*$. This implies that the maximand of $\Phi_1$ is greater than $d^*$, while the maximand of $\Phi_2$ is lower than $d^*$, and so the constraint that $d_1 \leq d_2$ is always binding in this case. In other words, the Agent should be given no discretion, and the Principal should simply follow policy $d^*$ in all states of the world.

Part (i) of this lemma states that if the Agent never wishes to choose a policy which is lower than the Principal's ideal policy $d(\theta)$, then the Principal's welfare from setting a lower bound is always decreasing in that bound, and therefore, this lower bound should be set as low as possible, and $d_1 = d_{\min}$. This is completely intuitive: if the only danger comes from the Agent choosing too high a policy, then the only tool the Principal should choose is to impose an upper bound on the allowed decisions. Similarly, if the only danger is from the Agent choosing too low a policy then only a lower bound should be imposed (part (ii) above). One example of a case with 'one-sided' errors might be the driving example discussed in the Introduction: if the government felt that the principal danger was of people driving too fast for the conditions — for instance, because they systematically ignored the negative externality of driving fast on the safety of other drivers — then there is no need to impose a minimum speed limit.

## 2.1 Comparative Statics

Suppose next that there is a change of Principal — what effect should this have on the optimal discretion interval? Suppose that the utility functions for the two Principals are $u(d,\theta)$ and $\hat{u}(d,\theta)$ respectively, and that the second has a preference for higher decisions $d$, all else equal, than the first, i.e.

$$\hat{u}_d(d,\theta) \geq u_d(d,\theta). \tag{8}$$

Let $\Phi_1$ and $\hat{\Phi}_1$ be their respective welfare levels from setting a lower bound, and let $\Phi_2$ and $\hat{\Phi}_2$ be their respective welfare levels from setting an upper bound. Then an examination of (7) above shows that

$$\begin{aligned} \hat{\Phi}_1'(d) &\geq \Phi_1'(d) \\ \hat{\Phi}_2'(d) &\geq \Phi_2'(d) \end{aligned}$$

and hence that the optimal lower bound is higher for the second Principal than the first, and similarly for the optimal upper bound. We summarise this in the following lemma:

**Lemma 3** *Suppose there are two Principals with utility functions $u$ and $\hat{u}$ which satisfy (8). If their optimal intervals of discretion are $[d_1, d_2]$ and $[\hat{d}_1, \hat{d}_2]$ respectively, then $\hat{d}_1 \geq d_1$ and $\hat{d}_2 \geq d_2$.*

The next result concerns the *tightness* of the limits to discretion rather than the levels of the bounds. Suppose we have two Agents (or populations of Agents) described by the stochastic response functions $G(d,\theta)$ and $\hat{G}(d,\theta)$

respectively. Let us say that the first Agent is *more reliable* than the second if

$$\begin{aligned} d &< d(\theta) \Rightarrow \hat{G}(d,\theta) \geq G(d,\theta) \\ d &> d(\theta) \Rightarrow \hat{G}(d,\theta) \leq G(d,\theta) \end{aligned} \tag{9}$$

for all $\theta$. In words, one Agent is more reliable than another if, given $\theta$, she is more likely to make a decision which is close to the Principal's ideal decision than the other Agent.

The next result is the main result of the paper, and states that if one Agent is more reliable than another the Principal should give that Agent more discretion in her decision-making, and also that the Principal obtains a higher expected utility from the more reliable Agent.

**Proposition 1** *Suppose there are two Agents (or two populations of Agents) whose decisions are described by the functions $G(d,\theta)$ and $\hat{G}(d,\theta)$ respectively. Let the Principal's optimal intervals for the two Agents be $[d_1, d_2]$ and $[\hat{d}_1, \hat{d}_2]$ respectively. Suppose that the first Agent is more reliable than the second in the sense of (9). Then the Principal's expected utility from delegating to the first Agent is higher than with the second, and*

$$[\hat{d}_1, \hat{d}_2] \subset [d_1, d_2] \ .$$

**Proof.** See Appendix.

Loosely speaking, this result states that increasing the risk of the Agent deviating from the Principal's preferred decision causes the Principal's utility to fall and causes the Agent to be given less discretion.[11]

A second attractive hypothesis is that increasing the variability of $\theta$, in some sense, should cause the Agent to be given greater discretion in order to respond to a more variable environment. For instance, if the distribution of $\theta$ was very tightly focused around $\theta^*$, say, then it is optimal for the Principal to offer very little discretion, and to set very tight bounds around the point $d(\theta^*)$ so as to attain approximately the first best. However, it appears to be difficult to obtain any straightforward comparative statics results on changing the distribution of $\theta$ without making strong assumptions about the functional forms of $G(d,\theta)$ and $u(d,\theta)$. For instance, from (7) it is clear that if we could assume that the function $G(d,\theta)u_d(d,\theta)$ was concave in $\theta$ for each $d$ then a mean-preserving spread of $\theta$ would reduce $\Phi_1'(d)$ at every point, and hence reduce the optimal lower bound.[12] Similarly, if $[1 - G(d,\theta)]u_d(d,\theta)$

[11] Of course, since her discretion is increased, any given agent is also made better off if the population of agents is perceived to be more reliable.

[12] See Rothschild and Stiglitz (1970).

was convex in $\theta$, then a mean-preserving spread of $\theta$ would cause the optimal upper bound to increase. However, without making further assumptions it is not possible to deduce that increasing the variability of the environment should cause greater discretion to be given to agents.

# 3 The Model with Divergent Tastes

In this section we try to provide a foundation for the rather 'reduced-form' character of the Agent's decision making used in the above analysis by modelling the preferences of the Agent in more detail.[13]

Suppose that both parties have utility functions of the form $u(\alpha, d, \theta)$ where $\alpha$ is a taste parameter which varies from person to person. Suppose, as before, that the marginal utility of increasing $d$, $u_d$, is increasing in $\theta$ and decreasing in $d$, but now suppose it is also increasing in $\alpha$. Therefore, a person's ideal decision, denoted $d(\alpha, \theta)$, is increasing in both $\alpha$ and $\theta$. Therefore, the Principal and Agent have *similar* preferences in the sense that each will choose a higher $d$ if a higher $\theta$ is observed, but the Agent's response function will shifted up or down compared to the Principal's, depending on whether the Agent's taste parameter is greater than or lower than the Principal's. A straightforward example of a suitable utility function is given by

$$u(\alpha, \theta, d) = -\frac{1}{2}(\theta + \alpha - d)^2 \tag{10}$$

in which case

$$d(\alpha, \theta) = \alpha + \theta \ .$$

Let the Principal's taste parameter be $\alpha_P$ and the agent's be $\alpha$. The Principal does not know $\alpha$ but holds a prior on $\alpha$ described by the distribution function $H(\alpha)$, so that $H(\alpha)$ is the probability that the Agent has a taste

[13] Another way of modelling the Agent's decision-making process, although from a theoretical point of view rather more *ad hoc* than the one considered in this section, is to suppose that the Agent is *incompetent*, rather than having divergent preferences. For instance, suppose that if both parties were omniscient they would each agree that the ideal policy given the state of the world $\theta$ is to choose $d(\theta)$. However, because of the possibly very complex problem involved in reaching a decision, having observed $\theta$ the Agent, if unconstrained, would wish to choose the policy $d(\theta) + \epsilon$, where $\epsilon$ is an error term due to bounded rationality, rather than $d(\theta)$. If the Principal foresees this danger, then in general it will be in her interest to limit the discretion of this Agent. Moreover, it is straightforward to obtain from Proposition 4 the result that the more likely $\epsilon$ is to be away from zero (i.e. the more incompetent is the agent), the less discretion she should be given.

parameter smaller than $\alpha$. The Principal believes that the variables $\alpha$ and $\theta$ are independently distributed.[14]

Given $d$ and $\theta$, let $\alpha(d, \theta)$ be that value of the taste parameter $\alpha$ for which the Agent with that taste parameter would wish to make decision $d$ if $\theta$ is observed, i.e. $d(\alpha(d, \theta), \theta) \equiv d$. Then $\alpha(d, \theta)$ is increasing in $d$ and decreasing in $\theta$. Given $\theta$, the probability that the Agent would wish to choose a policy $\delta \leq d$ is simply the probability that $\alpha$ is no greater than $\alpha(d, \theta)$, in other words

$$G(d, \theta) = H(\alpha(d, \theta)) \tag{11}$$

where $G$ is the Agent's stochastic response function used in the previous section.

In order to be able to use Proposition 1 above, we need to understand when one Agent is more or less reliable than another. So suppose we have two Agents (or populations of Agents) with distribution functions for their taste parameter given by $H(\alpha)$ and $\hat{H}(\alpha)$ respectively. Suppose further that

$$\begin{aligned} \hat{H}(\alpha) &\geq H(\alpha) \text{ if } \alpha \leq \alpha_P \\ \hat{H}(\alpha) &\leq H(\alpha) \text{ if } \alpha \geq \alpha_P \end{aligned} \tag{12}$$

so that the first Agent is more likely to have a taste parameter close to the Principal's than the other Agent. But if (12) holds, then so does (9). For suppose that $d \leq d(\alpha_P, \theta)$. Then by construction $\alpha(d, \theta) \leq \alpha_P$ and hence $\hat{H}(\alpha) \geq H(\alpha)$ from (12), and hence $\hat{G}(d, \theta) \geq G(d, \theta)$ from (11). (The other case where $d \geq d(\alpha_P, \theta)$ is similar.) Therefore, if the two distributions for $\alpha$ satisfy (12) then the first Agent is *more reliable* than the second.

Therefore, from Proposition 1 above, we may deduce the following result:

**Proposition 2** *Suppose there are two Agents (or populations of Agents) whose distribution functions for their taste parameters are $H(\alpha)$ and $\hat{H}(\alpha)$, these satisfying (12). Let the Principal's optimal intervals for the two Agents*

[14] It is worth mentioning a technical point concerning the relationship between this paper and Holmstrom (1984). Holmstrom has an agent with *known* preferences and with private information $y$ (which could be multi-dimensional). This paper models the agent has having private information $(\alpha, \theta)$, and known preferences given this information, but where the principal cares only about the state of the world $\theta$. At first glance, it might be imagined that the two models were the same provided that $(\alpha, \theta)$ was simply re-labelled as $y$. However, Holmstrom relies heavily in his analysis on an assumption that the principal and agent's preferences are *coherent*, and it is straightforward to show that if the principal is indifferent about one aspect of the agent's private information (which the agent does care about) then preferences cannot possibly be coherent. In this sense, the present model is a strict generalization of Holmstrom's model.

*be $[d_1, d_2]$ and $[\hat{d}_1, \hat{d}_2]$ respectively. Then the Principal's expected utility from delegating to the first Agent is higher than with the second, and*

$$[\hat{d}_1, \hat{d}_2] \subset [d_1, d_2].$$

# 4 Applications of the Model

Holmstrom (1984) proposes applications of his model which include a generalisation of Weitzman's (1974) price-versus-quantity dichotomy and limits to discretion in insurance contracts. In the following sections I discuss some further applications.

## 4.1 Mandatory Sentencing Policy

Consider some given class of crime (murder, for instance). A simple view of the purpose of a judge in a criminal trial is that he or she is there firstly to make sure that correct legal procedures take place, and secondly, and in the event of a guilty verdict, to determine the appropriate punishment given the particular circumstances of the case. A policy that in recent years has become more common is that of *mandatory* sentences for certain crimes.[15] This paper provides one rationale for this practice, namely, that judges may have divergent preferences from the political principal (e.g. Parliament or Congress).

We could represent the particular circumstances of a case by the parameter $\theta$, where $\theta$ could be a measure of the lack of 'mitigating circumstances' of the particular crime (which varies from case to case). For instance, if a women kills her drunken and violent husband who has been abusing her for years, then we might assign a low $\theta$ to the case. We could imagine, again over-simplifying somewhat, that a sentence serves two purposes: it is imposed as a kind of 'punishment' for the crime, and it acts as a 'deterrent' to potential future offenders. (We should also include the 'incapacitation' and 'rehabilitation' elements, but these are ignored here.) The punishment element in this decomposition would presumably decrease with the degree of mitigating circumstances of the particular case.

The 'agents' in this story are the various judges who might be involved in the case, and the 'principal' is the public body that is in overall control of

[15] Or rather, mandatory minimum punishments have become more common; maximum punishments for given crimes have almost always been stipulated. See, for instance, Robinson (1987) for a review of recent US policy in this area and for a discussion of the pros and cons of judicial discretion over sentencing policy.

sentencing policy (e.g. Parliament or Congress). Let the preference parameter $\alpha$ represent the weight placed on the punishment relative to the deterrent element by any particular judge, and let the decision variable $d$ represent the severity of the sentence. Let $u(\alpha, d, \theta)$ determine the preferences of a judge with parameter $\alpha$, and let $d(\alpha, \theta)$ denote the ideal punishment of the type-$\alpha$ judge for a case with circumstances $\theta$. One representation of preferences might be to set

$$u(\alpha, d, \theta) = u_1(d) + \alpha u_2(d, \theta)$$

where $u$ is decomposed into two parts: $u_1$ concerns the deterrent element of the punishment (perhaps not dependent on $\theta$), and $u_2$ is the punishment element (which does depend on the lack of mitigating circumstances $\theta$). With this formulation it makes sense to suppose that

$$u_2 > 0 \; ; \; \frac{\partial^2 u_2}{\partial d \partial \theta} \geq 0$$

in which case a smaller degree of mitigating circumstances leads all judges to impose a high punitive element of the sentence. The parameter $\alpha$, then, represents the 'punitiveness' of a given judge.

Suppose the principal has the preference parameter $\alpha_P$. The principal, if able to write a complete contract for judges specifying the punishment that should be imposed given the individual circumstances $\theta$, would instruct judges to impose the sentence $d(\alpha_P, \theta)$. However, in this context this seems quite impractical, and since it cannot monitor the circumstances $\theta$ of each case it must delegate the punishment decision to individual judges. The problem is that judges' views on reasonable sentences may differ from those of the principal, and some will be more lenient and some will be less lenient than it would wish. In order to counteract the danger of wayward judges, this paper suggests that the principal should impose mandatory sentencing on judges, and it should set a minimum sentence $d_1$ and a maximum sentence $d_2$ that any judge could impose for the given class of crime. Moreover, the more likely it is that judges have preferences that differ from those of the principal, the more restrictive should be these bounds (Proposition 2 above).[16,17]

[16]One reason why it is that mandatory maximum sentences are much more common than minimum sentences might be because it generally felt that the greatest danger comes from some judges being over-zealous in imposing punishments, rather than too timid, and so Lemma 2 could be relevant.

[17]An important consideration ignored in this discussion is that if a jury (in a criminal trial) knows that there is a mandatory minimum sentence that they consider excessive (e.g. death) they may be unwilling to convict even if they believe the suspect to be guilty.

## 4.2 Regulatory Discretion Over Pricing Policy

Consider a regulator deciding on pricing policy in a regulated industry. Specifically, suppose the regulated is setting a price cap for a near-monopoly and that this firm has an efficiency parameter $\theta$, where a high $\theta$ indicates a less efficient firm. Suppose that this parameter is known by the specialist regulatory agency but is not observed by the government. The decision variable $d$ represents the level of the price cap. A high value of $d$ benefits the near-monopoly as well as any entrants, but hurts consumers. Potential regulators differ in their opinion as to the appropriate weight to place on profits compared to the interests of consumers. Let $\alpha$ represent this weight and so

$$u(\alpha, d, \theta) = v(d) + \alpha\pi(d, \theta)$$

is the social welfare function of the type-$\alpha$ regulator, where $v(d)$ is consumer surplus with the price cap $d$ and $\pi(\theta, d)$ is industry profits with price cap $d$ and cost parameter $\theta$. Since a higher $\alpha$ represents a greater weight on firm profits, we would expect that $d(\alpha, \theta)$, the ideal choice of the type-$\alpha$ regulator, is increasing in $\alpha$ and $\theta$. This paper, then, suggests that if the government is unsure of the preferences of its regulator it should limit the discretion of the regulator by placing minimum and maximum bounds, $d_1$ and $d_2$ respectively, on the possible regimes the regulator may impose. The less likely it believes that the regulator's preferences will coincide with its own to be, the less discretion it should choose to give.[18]

## 4.3 Discretion Over Monetary Policy

Central bankers often have limits placed on their discretion in several areas. In particular, they are often strongly encouraged to stay within set bands for (i) certain measures of the money supply, and (ii) for the country's exchange rate as measured against other important currencies. It is probably true that maintaining policy credibility, in the face of time-inconsistency problems or of possible speculative attacks, is the dominant reason for such limits on central bank authority. However, the fact that the central banker may have preferences over monetary policy which differ from those of the govern-

[18] A real-world example of limited regulatory discretion concerns price cap regulation as faced by AT&T. This involves a requirement that the firm's average price (for a class of services) does not increase in percentage terms by more than a measure of the rate of inflation less a 'productivity' factor $X$. The magnitude of $X$ is chosen by the FCC, the industry regulator. However, the FCC faces limits on its choice of $X$ imposed by government (these limits being known as 'stabilizers'), and so must choose $X$ to lie within certain bounds $X_* \leq X \leq X^*$.

ment provides a second rationale for such practices. Consider, for instance, the tradeoff between employment and inflation discussed in Kydland and Prescott (1977). The decision $d$ could represent the tightness of the money supply (or alternatively, the exchange rate), and this is determined, within limits, by the central bank. The parameter $\theta$ could represent the various monetary and real shocks which affect the terms of the tradeoff between employment and inflation, and hence the ideal choice of $d$. The government and the central banker have preferences of the form $u(\alpha, \theta, d)$, where $\alpha$ could represent the 'conservativeness' of the preferences, i.e. a high value of $\alpha$ implies a greater weight being placed on reducing inflation at the expense of employment. A higher $\alpha$, then, implies a higher desired degree of monetary tightness $d$ (or a greater reluctance to devalue) for any given $\theta$. If $\alpha_P$ is the government's own preference parameter, then Proposition 2 implies that the more likely the central banker's preference parameter $\alpha$ is to be close to $\alpha_P$, the more discretion over monetary policy the bank should be given.[19]

# 5 Conclusions and Extensions

This paper has presented one way to think about the desirable limits to discretion when a principal delegates decisions to a well-informed agent with unknown preferences. If the principal could accurately observe all aspects of the environment and could write a complete contract specifying the policy to be taken by the agent in each circumstance, there is no benefit in limiting agency discretion. When this is not possible, however, she must allow the well-informed agent to decide policy within limits. The advantage of allowing full discretion is that agents have the freedom to respond to different environments in a desirable way. When there is a danger that the agent has policy preferences that differ from those of the principal, full discretion will allow such an agent the freedom to carry out undesirable policies. The correct degree of discretion must trade off these two effects. The main result of the paper (Proposition 2) was that the more likely an agent's preferences were to coincide with the principal's, the more discretion over policy the agent should be given.

The model contained a number of simplifying assumptions that could perhaps be abandoned in future work. These included the assumptions that the principal could only use the tool of limiting discretion to affect policy, that the agent's preferences were exogenous and could not be affected by interest group pressure, and that the decision variable $d$ was a scalar.

[19] See Delgado and Dumas (1992) for a discussion from a completely different perspective of the trade-offs between setting narrow and broad bands for exchange rate target zones.

It would be desirable to analyse the model without the first assumption. For instance, what would be the effect of introducing choice-contingent payments to the agent? However, I would regard this analysis as being complementary to that undertaken in this paper, rather than superseding it, because there are many settings in which choice-contingent payments do not occur. In addition, and as indicated earlier, I do not believe a full analysis of this case to be easily tractable. However, there are other mechanisms that are perhaps more practical which may also improve the situation. For instance, in a repeated principal-agent framework it is natural to suppose that the principal has the option of sacking the agent if she repeatedly makes unusual policy choices, and this may act to constrain agents with divergent preferences. Or it may be possible for the principal to employ more than one agent (perhaps at a cost) and then to take the 'average' policy recommendation. Provided that agents tastes $\alpha$ were roughly centered on the principal's parameter $\bar{\alpha}$ then taking the average would in many cases correspond to reducing the risk of divergent tastes which would then allow the principal to offer greater discretion.[20]

Secondly, I assumed that the agent's tastes were fixed. In the case of regulation, it may well be imagined that regulators are not simply *born* with given preferences over consumer versus industry interests, but that interest groups (e.g. an incumbent firm) can act to influence the preferences of the regulator, either by direct bribery or by more subtle inducements such as promises of lucrative future employment prospects. This has been the focus of the 'regulatory capture' literature discussed in the Introduction. These factors are of course likely to play an important role in shaping regulatory policy in some circumstances, but, in Britain at least, different regulators have clearly demonstrated different approaches to entry, industry profits and consumer welfare, and it is not plausible that this is due solely to differing interest group pressures. Rather, it seems likely that different regulators simply have different views, honourably held, on how to run a regulated industry. This paper presented the opposite polar approach to modelling preferences to that of the writers on 'capture' in the sense that regulatory preferences were assumed exogenous. Ideally one would wish to model regulatory behavior as a combination of the two approaches. For instance, we could think of $\alpha$, the relative weight the regulator places on industry profits compared to consumer welfare, as being composed of two factors: an exogenous taste parameter $\beta$ that was uncertain, and the magnitude of industry bribes $\gamma$, so that $\alpha = \beta + \phi(\gamma)$ for some increasing function $\phi(\gamma)$.

[20] This is one argument for using a committee system for deciding monetary policy in central banks, as is the case, for instance, in the German Bundesbank.

Finally there is the more technical point about the decision variable $d$ being a scalar. I do not believe that a generalisation to multi-dimensional decisions would be straightforward. For one thing, the choice of the set of possible control sets will be problematic. (In the scalar case, the set of choice sets that I considered was the set of intervals $[d_1, d_2]$.) Just as in the scalar case, the set of all compact sets is too difficult a set to deal with (I believe), but unlike the scalar case even the set of all convex sets, for instance, will again be too broad. Instead, in order to gain tractable results it may be that *ad hoc* families of sets such as rectangles or circles would need to be considered, and that because of the artificial nature of this assumption, simple results connecting the dispersion of tastes and the degree of discretion could be difficult to obtain. Moreover, in a multi-dimensional setting it will often be precisely the *shape* of the choice set that is of interest. For instance, consider the problem of finding the optimal form of price cap regulation for a multiproduct firm.[21] We can think of this problem within the present framework as one of delegating the decision over pricing policy to a well-informed agent — the profit-maximising firm — who has preferences that diverge from social welfare. In this case, what is of interest is the shape of the allowed set of prices (e.g. to what extent does it resemble such commonly used mechanisms as 'average revenue' price cap regulation) as much as anything else.

# A Proof of Lemma 1

(i) Fix $\bar{d} > d^*$. Since $U(\cdot)$ is concave and is maximized at $d = d^*$ we know that $U'(\bar{d}) \leq 0$, i.e. $E_\theta u_d(\bar{d}, \theta) \leq 0$. Suppose $\bar{\theta}$ is such that $d(\bar{\theta}) = \bar{d}$. (If no such $\bar{\theta}$ can be found then $\bar{d}$ lies outside the range of the Principal's ideal policies, and the following argument can easily be adapted to cope with this degenerate case.) In particular, since $u_d$ is increasing in $\theta$ and $u_d(\bar{d}, \bar{\theta}) = 0$, it follows that $u_d(\bar{d}, \theta) \leq 0$ for $\theta \leq \bar{\theta}$ and $u_d(\bar{d}, \theta) \geq 0$ for $\theta \geq \bar{\theta}$. Then from (7) we obtain

$$\begin{aligned}
\Phi_1'(\bar{d}) &= \int_{\theta_{\min}}^{\bar{\theta}} G(\bar{d}, \theta) u_d(\bar{d}, \theta) f(\theta)\, d\theta + \int_{\bar{\theta}}^{\theta_{\max}} G(\bar{d}, \theta) u_d(\bar{d}, \theta) f(\theta)\, d\theta \\
&\leq \int_{\theta_{\min}}^{\bar{\theta}} G(\bar{d}, \bar{\theta}) u_d(\bar{d}, \theta) f(\theta)\, d\theta + \int_{\bar{\theta}}^{\theta_{\max}} G(\bar{d}, \bar{\theta}) u_d(\bar{d}, \theta) f(\theta)\, d\theta \\
&= G(\bar{d}, \bar{\theta}) \times E_\theta u_d(\bar{d}, \theta) \\
&\leq 0
\end{aligned}$$

[21] For a discussion of this point, see Armstrong et al. (1994, section 3.3.2).

where the first inequality follows from (1).

The proof of (ii) Similar to part (i). •

# B Proof of Proposition 1

Let $(\Phi_1(d), \Phi_2(d))$ and $(\hat{\Phi}_1(d), \hat{\Phi}_2(d))$ be the two pairs of payoff functions for setting upper and lower bounds for the Principal that result from the two Agents. Then the following argument is similar to that used in the proof of Lemma 1. Fix a given $\bar{d}$ and let $\bar{\theta}$ be given by $d(\bar{\theta}) = \bar{d}$. Then because $u$ is concave in $d$ and $d(\theta)$ is increasing in $\theta$, it follows that $\bar{d} \geq d(\theta)$ if $\theta \leq \bar{\theta}$ and $\bar{d} \leq d(\theta)$ otherwise, and also that $u_d(\bar{d}, \theta) \leq 0$ if $\theta \leq \bar{\theta}$ and $u_d(\bar{d}, \theta) \geq 0$ otherwise. Putting these inequalities together and using (7) we obtain

$$\begin{aligned} \Phi_1'(\bar{d}) &= \int_{\theta_{\min}}^{\bar{\theta}} G(\bar{d}, \theta) u_d(\bar{d}, \theta) f(\theta)\, d\theta + \int_{\bar{\theta}}^{\theta_{\max}} G(\bar{d}, \theta) u_d(\bar{d}, \theta) f(\theta)\, d\theta \quad (13) \\ &\leq \int_{\theta_{\min}}^{\bar{\theta}} \hat{G}(\bar{d}, \bar{\theta}) u_d(\bar{d}, \theta) f(\theta)\, d\theta + \int_{\bar{\theta}}^{\theta_{\max}} \hat{G}(\bar{d}, \bar{\theta}) u_d(\bar{d}, \theta) f(\theta)\, d\theta \\ &= \hat{\Phi}_1'(\bar{d}) \end{aligned}$$

where the inequality follows from (9). A similar argument shows that

$$\Phi_2'(d) \geq \hat{\Phi}_2'(d) \ . \tag{14}$$

Inequality (13) shows that the Principal will necessarily set a higher lower bound for the less reliable Agent, while (14) shows that she will set a lower upper bound for the less reliable Agent. In other words, the more reliable agent will be offered a greater degree of discretion.

It is also intuitive that the Principal will obtain greater expected utility from a more reliable agent. To see formally that this is true, we can integrate the third term in (3) by parts to obtain

$$w(d_1, d_2, \theta) = u(d_2, \theta) - \int_{d_1}^{d_2} u_d(\delta, \theta) G(\delta, \theta)\, d\delta \ . \tag{15}$$

Let $w$ and $\hat{w}$ denote the two payoff functions given $\theta$ resulting from the two Agents. But it is straightforward to check that if the second Agent is less reliable than the first, then $w(d_1, d_2, \theta) \geq \hat{w}(d_1, d_2, \theta)$ for all $(d_1, d_2, \theta)$. For instance, suppose $d_1 \leq d(\theta) \leq d_2$. Then from (15) we obtain

$$\begin{aligned} w(d_1, d_2, \theta) &= u(d_2, \theta) - \int_{d_1}^{d(\theta)} u_d(\delta, \theta) G(\delta, \theta) d\delta - \int_{d(\theta)}^{d_2} u_d(\delta, \theta) G(\delta, \theta) d\delta \\ &\geq u(d_2, \theta) - \int_{d_1}^{d(\theta)} u_d(\delta, \theta) \hat{G}(\delta, \theta) d\delta - \int_{d(\theta)}^{d_2} u_d(\delta, \theta) \hat{G}(\delta, \theta) d\delta \\ &= \hat{w}(d_1, d_2, \theta) \end{aligned}$$

where the inequality follows from (9) and the fact that $u_d(\delta, \theta) \geq 0$ if $\delta \leq d(\theta)$ and $u_d(\delta, \theta) \leq 0$ if $\delta \geq d(\theta)$. (The case where $d(\theta)$ lies outside the interval $[d_1, d_2]$ is similar.) Therefore $W(d_1, d_2) \geq \hat{W}(d_1, d_2)$, where $W$ and $\hat{W}$ are the expected welfare functions resulting from the two Agents. •

## REFERENCES


Aghion, P. and J. Tirole (1994). *Formal and Real Authority in Organizations.* Working Paper no. 37, Institut d'Economie Industrielle, Toulouse.

Armstrong, M. (1994). *Delegation and Discretion.* Discussion paper no. 9421, Department of Economics, University of Southampton.

Armstrong, M., Cowan, S. and Vickers, J. (1994). *Regulatory Reform — Economic Analysis and British Experience.* Cambridge: MIT Press.

Athey, S., Gans, J., Schaefer, S. and Stern, S. (1994). *The Allocation of Decisions in Organizations.* Mimeo, Stanford University.

De Bijl, P. (1995). *Delegation of Responsibility in Organizations.* Mimeo, Tilburg University.

Delgado, F. and B. Dumas (1992). Target Zones, Broad and Narrow. In P. Krugman and M. Miller, (eds), *Exchange Rate Targets and Currency Bands.* Cambridge: Cambridge University Press.

Gilbert, R. and Newbery, D. (1994). The Dynamic Efficiency of Regulatory Constitutions. *Rand Journal of Economics* 25: 538–554.

Greenwald, B.C. (1984). Rate Base Selection and the Structure of Regulation. *Rand Journal of Economics* 15: 85-95.

Holmstrom, B. (1984). On the Theory of Delegation. In M. Boyer and R.E. Kihlstrom (eds) *Bayesian Models in Economic Theory.* Elsevier Science Publishers.

Holmstrom, B. and P. Milgrom (1991). Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. *Journal of Law, Economics and Organization* 7: 24– 56.

Kydland, F. and Prescott, E. (1977). Rules rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy* 85: 473-492.

Laffont, J.–J. and Tirole, J. (1993). *A Theory of Incentives in Procurement and Regulation.* Cambridge: MIT Press.

Peltzman, S. (1976). Towards a More General Theory of Regulation. *Journal of Law and Economics* 14: 109–148.

Posner, R. (1974). Theories of Economic Regulation. *Bell Journal of Economics* 5: 335–358.

Rogerson, P. (1987). A Sentencing System for the 21st Century? *Texas Law Review* 66: 1–61.

Rogoff, K. (1985). The Optimal Degree of Commitment to an Intermediate Monetary Target. *Quarterly Journal of Economics* 100: 1196–1190.

Rothschild, M. and Stiglitz, J. (1970). Increasing Risk: I. A Definition. *Journal of Economic Theory* 2: 225–243.

Stigler, G. (1971). The Theory of Economic Regulation. *Bell Journal of Economics* 2: 3–21.

Vickers, J. (1985). Delegation and the Theory of the Firm. *Economic Journal,* Conference Papers 95: 138-147.
Weitzman, M. (1974). Prices vs. Quantities. *Review of Economic Studies* 41: 477–491.