A model of delegated project choice

Armstrong, Mark and Vickers, John

University College London (UCL)

18 March 2009

A Model of Delegated Project Choice*

Mark Armstrong
Department of Economics
University College London

John Vickers
All Souls College
University of Oxford

March 2009

Abstract

We present a model in which a principal delegates the choice of project to an agent with different preferences. The principal determines the set of projects from which the agent may choose. The principal can verify the characteristics of the project chosen by the agent, but does not know which other projects were available to the agent. We consider situations where the collection of available projects is exogenous to the agent but uncertain, where the principal cannot observe the agent's payoff, where the principal can pay the agent to choose a desirable project, and where the agent must invest effort to discover a project.

Keywords: Delegation, principal-agent, rules, merger policy, hurdle rates.

1 Introduction

In the main model in this paper we present an analysis of a principal-agent problem in which the principal can influence the agent's behaviour not by outcome-contingent rewards but by specifying what the agent is, and is not, allowed to do. The agent, whose preferences differ from those of the principal, will select from her available projects the permitted project that best serves her interests. The principal can verify whether or not the selected project is indeed within the permitted set, but cannot observe the number or characteristics of the projects available to the agent. How then should the principal specify the set of projects from which the agent can choose?

*We are grateful to Daron Acemoglu, V. Bhaskar, Martin Cripps, James Dow, Florian Englmaier, Bengt Holmstrom, Michael Katz, Niko Matouschek, Meg Meyer, Tima Mylovanov, Barry Nalebuff, Eric Rasmusen, David Sappington, Karl Schlag, Aggey Semenov, Joel Sobel, Glen Weyl, Jidong Zhou, and to three referees for comments and discussion. Armstrong gratefully acknowledges the support of the Economic and Social Research Council (UK).

One application of our analysis is to an important question in competition policy: should mergers be permitted provided that they are expected not to diminish "total welfare", or should the policy standard be in terms of consumer welfare? Even if the ultimate policy objective is total welfare, the answer is not obvious because the merger proposals brought forward by (profit-seeking) firms are selected by them from a wider pool of potential mergers. As we shall see, because of this selection effect, it could be that total welfare is higher on average under a consumer standard than a total welfare standard.

Another example of constrained delegation relates to decision making within a firm. The owners of a company may, within limits, delegate investment project choice to a manager, where the owners are interested in profits, while the manager enjoys private benefits from projects and/or private costs from searching for projects. We show that, in such a setting, it will not generally be optimal for the owners to approve all projects that would be profitable. A quite different, behavioural, application of our model is to a single agent whose future tastes may differ from her current tastes. This agent may wish to constrain her future choices now, if feasible, in order to maximize her well-being over time.

More generally, our analysis addresses an aspect of the theory of optimal *rules*—namely, the relationship between the ultimate objective of the rule-setter and the optimal rule to commit to. That relationship is not straightforward inasmuch as the likely consequences of a rule—including for the attainment of the ultimate objective—depend on the responses of agents seeking to maximize, within the rules, their own objectives. The interplay between rules and the responses that they induce is at the heart of our analysis.

Our goal is not to compare alternative *ad hoc* rules, such as the consumer and total welfare standards in the merger policy example. Rather, it is to characterize optimal rules in terms of the fundamentals of our models. Sometimes optimal rules are found to have strikingly simple forms. Our benchmark model in section 3 analyzes a setting in which the agent chooses one project from an exogenous, but uncertain, number of available projects. Contingent monetary incentives are ruled out, and the principal optimally restricts agent choice in a way that forbids some projects that are moderately good, in the uncertain hope of inducing the agent instead to choose a project that is better for the principal. This bias is akin to putting less weight on the agent's payoff than is in the true welfare function.

In section 4 we present three variants of this benchmark model. Situations in which the principal cannot observe the agent's payoffs are discussed in section 4.1. This variant might apply to project choice within an organization, where the principal sets a (possibly high) hurdle rate for the profitability of permitted projects. Second, in section 4.2 we discuss the impact of monetary incentives to choose a good project. When the agent is liquidity constrained, it may be preferable to restrict the agent's freedom to choose projects than to pay the agent to choose a good project. Finally, in section 4.3 we analyze a setting where the agent influences the likelihood of finding a project by exerting costly effort. Here, we show that the principal optimally sets a *linear* permission rule. In addition, to induce greater effort, the principal allows some projects that are detrimental for his interests; this bias is akin to putting more weight on the agent's payoff than is in his true welfare function.

Some other papers have examined situations in which a principal delegates decision-making to a (potentially) better-informed agent whose preferences differ from those of the principal, and where contingent transfers between principal and agent are ruled out. Aghion and Tirole (1997) show how, depending on information structure and payoff alignment, it may be optimal for a principal to delegate full decision-making power to a potentially better-informed agent. The principal's loss of control over project choice can be outweighed by advantages in terms of encouraging the agent initiative to develop and gather information about projects. In like vein Baker, Gibbons, and Murphy (1999), though they deny formal delegation of authority, examine informal delegation through repeated-game relational contracts. Even an informed principal able to observe project payoffs may refrain from vetoing ones that yield him poor payoffs in order to promote search incentives for the agent.

Our work is closer to the models which analyze *constrained* delegation, where the agent can make decisions but only within specified limits and the principal's problem is to decide how much leeway to give the agent. This literature was initiated by Holmstrom (1984), and the elements of his model go as follows. There is a set of projects, indexed by a scalar variable d which takes values in some large interval D, of which one needs to be chosen. A given project generates payoffs to the two parties which depend on the state of the world, represented by θ, and only the agent observes this parameter. The preferences of the principal and agent may differ, and if project d is chosen when the state is θ the principal obtains payoff $Y_P(d,\theta)$ and the agent has payoff $Y_A(d,\theta)$. The principal's problem is to choose a permission set, say $\mathcal{D} \subset D$, in order to maximize his expected payoff (given his prior on the distribution of θ), given that the agent will choose her most preferred project in $\mathcal{D}$ given the state θ. This "delegation problem" coincides with the "mechanism design problem" where the agent makes an announcement about the claimed state of the world, $\hat{\theta}$, and the principal commits to a decision rule $d(\hat{\theta})$ which maps the announcement to the implemented project. The two approaches are equivalent since the principal never directly observes the true θ and by making a suitable announcement $\hat{\theta}$ the agent can implement any project in the range of the rule $d(\cdot)$.

Holmstrom mostly limits attention to the situation in which the permission set $\mathcal{D}$ is an interval. Subject to this assumption (and other regularity conditions), he shows that an agent whose preferences are closer to the principal's will be given wider discretion. (This result has subsequently sometimes been termed the "ally principle".) Following Holmstrom's initial contribution, subsequent papers have analyzed when interval delegation is in fact the optimal delegation policy, making the additional assumption that θ is a scalar variable.[1] Melumad and Shibano (1991) were the first to calculate optimal permission sets, in the special case where preferences were quadratic and where θ was uniformly distributed. They found that interval delegation was optimal when principal and agent have ideal projects which are similarly responsive to the state θ, but that otherwise it could be optimal to have "holes" in $\mathcal{D}$. Martimort and Semenov (2006) find a sufficient condition on the distribution of θ for interval delegation to be optimal. Alonso and Matouschek (2008) systematically investigate when interval delegation is optimal, and they generalize Melumad and Shibano's insight that the relative responsiveness of preferred decisions to the state is the key factor for this. They show that when interval delegation is sub-optimal the ally

[1] Szalay (2005) presents an interesting variant on this delegation problem in which interval delegation is often sub-optimal. In his model, there is no divergence in preferences between the principal and agent, but the agent incurs a private cost to observe θ. He shows that it can be optimal for the principal to remove intermediate projects from $\mathcal{D}$ so that the agent is forced to choose between relatively extreme options, and this sharpens the agent's incentive to discover θ. The benefit from this can outweigh the principal's loss due to the possibility that the ideal project given the realized state is forbidden.

principle need not hold and an agent with preferences more aligned with those of the principal might optimally be given less discretion.[2]

Those models in the Holmstrom tradition differ from ours in respect of project specification and the form of asymmetric information. In particular, they characterize each project by a scalar parameter, all projects are always feasible, and the agent has private information about a payoff-relevant state of the world. In our model, by contrast, payoffs of the chosen project to both principal and agent are known, but only a finite collection of projects is feasible and only the agent knows what those projects are. Like the papers discussed above, our aim is to characterize the optimal delegation set from which the agent chooses, but (except for section 4.1) in the two-dimensional setting where the principal can observe both his own and the agent's payoff from the project chosen by the agent.[3]

2 Motivations for the Analysis

2.1 Welfare standards in merger policy

An important debate in antitrust policy concerns the appropriate welfare standard to use when deciding whether to prohibit a merger (or some other form of conduct).[4] The two leading contenders are a *total welfare* standard, where mergers are evaluated according to whether they decrease the unweighted sum of producer and consumer surplus, and a *consumer welfare* standard, where mergers detrimental to consumers are blocked. Many economic commentators feel that antitrust policy should aim to

[2] Kovác and Mylovanov (2007) discuss the impact of stochastic delegation schemes, in which the agent chooses from a constrained set of *lotteries* over projects rather than from a subset of projects themselves. They show, broadly speaking, that when interval delegation is optimal in Alonso and Matouschek's deterministic framework no extra benefit is generated by stochastic schemes, but otherwise the principal can sometimes do better by offering lotteries over projects.

[3] A recent paper which also investigates a two-dimensional delegation problem is Amador, Werning, and Angeletos (2006). There, an agent with quasi-hyperbolic preferences has wealth which she consumes over two periods. If there were no uncertainty about her preferences, she would gain by committing to a fixed consumption path at time zero. However, she will receive a utility shock in period 1 and this uncertainty gives a motive allow some flexibility in consumption. Amador *et al.* show that the time-zero agent limits her consumption possibilities to a subset of feasible consumption paths, and they find a condition which implies that the optimal delegation set simply involves placing a ceiling on first-period consumption.

[4] In practice, many mergers giving rise to competition concerns are permitted subject to conditions (e.g., asset disposals) rather than prohibited altogether. Such merger remedies are beyond the scope of the illustrative example discussed here.

maximize total welfare, whereas in most jurisdictions the focus is more on consumer welfare alone. See Farrell and Katz (2006) for an overview of the issues.

One purpose of this paper is to examine a particular strategic reason, discussed by Lyons (2002) and Fridolfsson (2007), to depart from the regulator's true welfare standard, which is that a firm may have a *choice* of merger possibilities. A less profitable merger might be better for total welfare, but will not be chosen under a total welfare standard. To illustrate, consider Figure 1, which is similar to those presented in section IV.B in Farrell and Katz (2006).

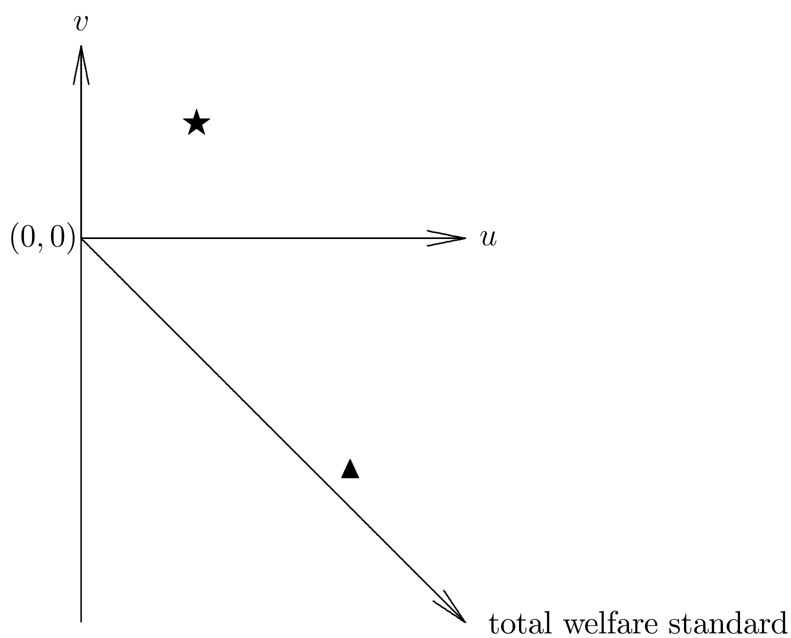

Figure 1: The Impact of Welfare Standard on Chosen Mergers

Here, u represents the gain in total profit resulting from a merger, while v measures the resulting net gain (which may be negative) to consumers. Suppose that u and v are verifiable once a merger is proposed to the competition authority. If the regulator follows a total welfare standard, he will permit any merger which lies above the negatively-sloped line in the figure. Suppose the firm has two mergers to choose from, depicted by ▲ and ★ on the figure. With a total welfare standard, the firm will choose the merger with the higher u payoff, i.e., the ▲ merger. However, the regulator would prefer the alternative ★ since that yields higher total welfare. If the regulator instead imposed a consumer welfare standard, so that only those mergers which lie above the horizontal line $v = 0$ are permitted, then the firm will be forced to choose the preferred merger. In this case, a regulator wishing to maximize total welfare is better off if he imposes a consumer welfare standard. As Farrell and Katz (2006, page 17) put it: "if we want to maximize gains in total surplus (northeasterly movements as shown in Figure [1]) and firms always push eastwards, there is something to be said for someone adding a northerly force." Nevertheless, there is a potential cost to adopting a consumer welfare standard: if the ▲ merger turns out to be the *only* possible merger then a consumer welfare standard will not permit this even though the merger will improve total welfare. Thus, the choice of welfare standard will depend on the number of possible mergers and the distribution of profit and consumer surplus gains for a possible merger.[5]

2.2 Project choice within an organization

Within a firm, a manager may be given authority to select projects to implement, subject to limits imposed by the firm's owners. For instance, a particular project may yield an expected profit to the shareholders of v, together with a personal benefit to the manager of u. In some situations, u might be observable to the owners. For example, the manager might enjoy, or dislike, implementing "large" projects (e.g., as measured by the initial investment needed), and the size of the project could be measured by the principal. In such cases, the owners could permit the manager to implement any project that lies in some subset $\mathcal{D}$ of possible realizations of (u, v).

In other cases, u might be a private benefit for the manager, unobserved by the principal. Here, the owners can condition permission only on the value of v, not u, and only those projects with profitability $v \geq R$, say, might be permitted. The threshold R is sometimes termed a "hurdle rate". If the manager sometimes has a choice of projects to implement, the principal may wish to set a hurdle rate which strictly exceeds the firm's own cost of capital, so that some projects which have positive value to shareholders are not permitted. The reason is similar to the why the competition authority might wish to impose a consumer-oriented welfare standard: by forbidding mediocre projects, the manager will be forced to choose her

[5] A slightly different interpretation is that an industry has a choice of merger possibilities. For instance, a particular firm may be a target for takeover in an industry and have several potential bidders to take it over. Each of these potential takeovers generates its own impact on consumer surplus and profit. A plausible situation is that winning bidder will be the firm which generates the highest gain in profit subject to meeting the required welfare standard.

next best permitted project, which could be significantly profitable. Although it is often feasible to offer contingent monetary incentives to induce a manager to choose a profitable project, in section 4.2 we show how it can be better for the owners to ban mediocre projects than to reward the manager to choose a very profitable project.

There is a rich literature which examines why, and the extent to which, firms set high hurdle rates for project approval. Much of the theoretical work suggests that high hurdle rates can be used by a firm's owners to reduce a manager's rent in the context where the manager has private information about a project's profitability. (See Lambert, 2001, section 4, for a survey of this literature.) In such models, the manager has only one project he can propose. In an extension to the models in that tradition, Lambert (2001, section 4.3) analyzes a model where the manager obtains non-pecuniary benefits from implementing a project of given size, and he shows that if the manager is risk-averse and her benefits are increasing in the size of the project, it is optimal to implement a hurdle rate above the cost-of-capital.

Closer to our approach is Dow and Raposo (2005), where a manager potentially has two different types of project available, one of which is likely to be more profitable but also more costly to implement. As in our model, the firm's owners cannot observe how many projects are available, and the manager proposes only her preferred project whenever she has a choice of project. The focus of that paper is on the compensation schemes offered to the manager, rather than restricting the set of permitted projects, although in section IIB they show how compensation can be structured so that one project is effectively banned: as in our model this has the benefit that shareholders obtain the project they prefer when there is a choice of project, but it has the cost when the less good project is the only project available that a moderately profitable opportunity is forgone.

3 Benchmark Model: Choosing a Project

A *principal* ("he") delegates the choice of project to an *agent* ("she"). There may be several projects for the agent to choose from, although only one can be implemented over the relevant time horizon. A *project* is fully described by two parameters, u and v. The agent's payoff if the type-(u, v) project is implemented is u, while the payoff to the principal, who is assumed to be risk-neutral, is $v + \alpha u$. (If no project is implemented, each party obtains payoff of zero.) Here, $\alpha \geq 0$ represents the weight the principal places on the agent's interests, and v represents factors specific to the principal's interests. The parameter α might reflect a true regard for the agent's payoff (as in the merger application when profits carry some weight in social welfare), and/or it might reflect a trade-off between allowing the agent wider project choice—and so a greater chance of on-the-job benefits u—and paying her a higher salary.[6] For example, if the principal jointly chooses the permitted set of projects and the salary level to meet the participation constraint of a risk-neutral agent, and u is measured in money terms, then the following analysis applies with $\alpha = 1$.[7]

Each project is an independent draw from the same distribution for (u, v). Since the agent will never propose a project with a negative payoff, without loss of generality we suppose that only non-negative u are realized. The marginal density of $u \geq 0$ is $f(u)$. The conditional density of v given u is denoted $g(v, u)$ and the associated conditional distribution function is $G(v, u)$. Here, v can be positive or negative. Suppose that the support of (u, v) is a rectangle $[0, u_{\max}] \times [v_{\min}, v_{\max}]$, where $v_{\min} \leq 0 \leq v_{\max}$ so that $(0, 0)$ lies in the support of (u, v). Finally, suppose that both f and g are continuously differentiable and non-zero on the support of (u, v).

In this benchmark model, the number of projects is random and the probability that the agent has exactly $n \geq 0$ available projects is q_n. (Our analysis covers the case where there are surely N projects, but the analysis is no easier for that case. Indeed, we will see that n being a Poisson variable is the easiest example to analyze.) Suppose that the project characteristics (u, v) are distributed independently of n.

The principal delegates the choice of project to the agent. We assume in this benchmark model that it is impossible, or not credible, for the principal to give contingent monetary incentives to the agent to choose a desirable project. Once the agent selects a particular project from her menu of possible projects, we assume that the characteristics of the chosen project—and *only* that project—are verifiable. This assumption may be appropriate if the principal has substantial costs associated with auditing a given project's characteristics, and/or the agent has significant costs

[6] Aghion and Tirole (1997) refer to this benefit of giving the agent freedom to choose projects as the "participation" benefit of delegation.

[7] At the other extreme, if the agent is infinitely risk-averse and cares only about her minimum income over all outcomes, then the following analysis applies when $\alpha = 0$. Such an agent is unwilling to trade-off her salary against the uncertain prospect of on-the-job benefits.

associated with preparing a credible project proposal. (Each of these is arguably the case in the merger scenario, for example.)

Recall that in the Holmstrom (1984) setting the "delegation problem" coincided with the "mechanism design problem" in which the agent reports her (scalar) private information. In our setting her private information consists of the number and nature of available projects, and the mechanism design approach would involve the principal committing to a rule that determines which project is chosen as a function of the agent's announced list of available projects.[8] Under our assumption that only the chosen project's characteristics are verifiable, the (deterministic) mechanism design problem boils down to the delegation problem of choosing the set of permitted projects. To see this, note that the set of projects can be partitioned into two subsets: the set of projects, say $\mathcal{D}$, which, by making suitable reports of other projects, could be chosen for implementation under the principal's decision rule, and those projects which are never implemented by the principal's rule. Faced with this rule, the agent will simply choose her preferred available project (if any) in the former set, and announce any other projects required to implement that choice. Clearly, this mechanism is equivalent to the delegation problem where the agent can directly choose any project in the set $\mathcal{D}$.

In sum, the mechanism we analyze is as follows: before the agent has any private information, the principal commits to a (measurable) "permission set" (or delegation set) of projects, denoted $\mathcal{D} \subset [0, u_{\max}] \times [v_{\min}, v_{\max}]$, and the agent can then implement any project in $\mathcal{D}$.[9,10] (See the concluding section for discussion of more elaborate

[8]If the project selected under the principal's rule turns out not to have the agent's claimed characteristics, suppose the principal does not implement any project.

[9]It is important to emphasize the assumption here, as in delegation problems more generally, of commitment. We will see that the principal excludes some *ex post* desirable projects from the permission set. But in the absence of commitment power, if the agent reports that she has no permitted project, the principal may have an incentive to "renegotiate", and to widen the set of permitted projects. In section 4.3 we find the opposite bias, in which the principal permits some projects which are undesirable to him *ex post*, and this again requires commitment power to be credible. See Baker, Gibbons, and Murphy (1999) and Alonso and Matouschek (2007) for analyses of how commitment power can be endogenously generated with repeated interaction.

[10]We restrict attention to deterministic mechanisms, mainly because it is hard to imagine being able to commit to or implement a stochastic mechanism in practice. It is possible that commitment to a stochastic scheme, if feasible, could do better for the principal than a deterministic scheme. For example, suppose that $\alpha = 0$, that $n = 2$ for sure, and that $(u, v) = (0.5, 1)$ with probability 0.5 and $(u, v) = (0.9, 0.1)$ with the same probability. In this case, the optimal deterministic scheme only permits the principal's favoured project, $(0.5, 1)$. However, permitting the project $(0.9, 0.1)$ with probability 0.5 yields the principal a higher expected payoff than banning it altogether: in

delegation schemes.)

For each u let $\mathcal{D}_u = \{v \text{ such that } (u,v) \in \mathcal{D}\}$ be the set of u-projects which are permitted, and let

$$p(u) = \int_{v \in \mathcal{D}_u} g(v,u) dv$$

be the proportion of type-u projects which are permitted. Let

$$x(u) = 1 - \int_u^{u_{\max}} p(z) f(z) \, dz$$

to be the probability that any given project either has agent payoff z less than u or is not permitted. Note that $x(\cdot)$ is continuous, and when it is differentiable its derivative is

$$x'(u) = p(u) f(u) \; . \tag{1}$$

(Since $x(\cdot)$ is weakly increasing it is differentiable almost everywhere.) If there are n available projects, the probability that each project is either banned or generates agent payoff less than u is $(x(u))^n$. Summing over n implies that the probability that each available project is either banned or generates agent payoff less than u is $\phi(x(u))$, where $\phi(x) \equiv \sum_{n=0}^{\infty} q_n x^n$ is the *probability generating function* (or PGF) associated with the random variable n. It follows that the density of the agent's preferred permitted project (where this exists) is $\frac{d}{du}\phi(x(u))$. Useful properties of PGFs which we will sometimes use are that they are well-defined on the relevant interval $0 \leq x \leq 1$ and smooth, convex and increasing over this interval.

The principal's payoff with the permission set $\mathcal{D}$ is therefore

$$\int_0^{u_{\max}} \{E[v \mid (u,v) \in \mathcal{D}] + \alpha u\} \frac{d}{du}\phi(x(u)) \, du$$

$$= \int_0^{u_{\max}} \left\{ \int_{v \in \mathcal{D}_u} v g(v,u) dv + \alpha u p(u) \right\} f(u) \phi'(x(u)) \, du \; . \tag{2}$$

The principal's problem is to maximize expression (2) taking into account the relationship between and p and x in (1) and the endpoint constraint $x(u_{\max}) = 1$. The following lemma shows that the optimal permission set takes a "threshold" form:

Lemma 1 *In the optimal policy there exists a threshold rule $r(\cdot)$ such that*

$$(u,v) \in \mathcal{D} \textit{ if and only if } v \geq r(u) \; .$$

both cases the agent would choose the principal's preferred project if she could, but if that is not available the stochastic scheme would still allow some chance of a desirable project being chosen.

Proof. From (1), the function $x(\cdot)$ depends on $\mathcal{D}$ only via the "sufficient statistic" $p(u)$, not on the particular v-projects which are permitted given u. Therefore, for any candidate function $p(u)$ the principal might as well permit those particular v-projects which maximize the term $\{\cdot\}$ in the (2), subject to the constraint that the proportion of type-u projects is $p(u)$. But the problem of choosing the set $\mathcal{D}_u$ in order to

$$\text{maximize} \int_{v\in\mathcal{D}_u} vg(v,u)dv \text{ subject to } \int_{v\in\mathcal{D}_u} g(v,u)dv = p(u)$$

is solved by permitting the projects with the highest v so that the proportion of permitted projects is $p(u)$, i.e., that $\mathcal{D}_u = \{v \text{ such that } v \geq r(u)\}$ for some $r(u)$. ■

Thus, the problem simplifies to the choice of threshold rule $r(\cdot)$ rather than the choice of permission set $\mathcal{D}$. (A similar argument is valid in the search model in section 4.3, although not for the private benefits model in section 4.1.) Figure 2 depicts a possible threshold rule $r(\cdot)$, and shows $x(u)$ as the measure of the shaded area.

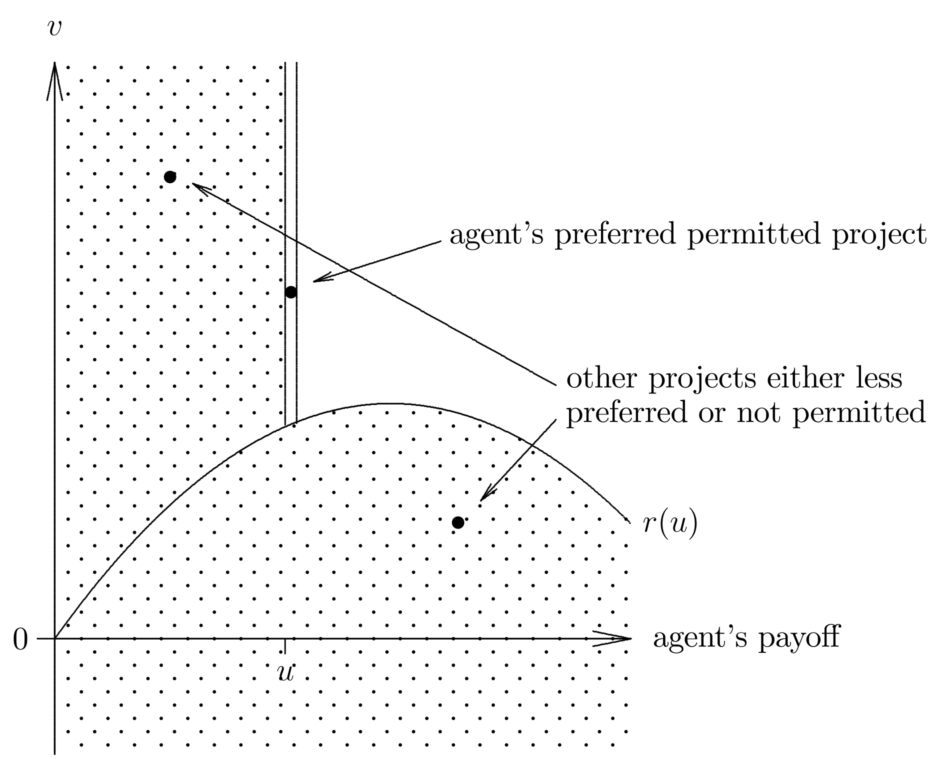

Figure 2: The Agent's Preferred Permitted Project

Define $V(r,u)$ to be the expected value of v given that the project has agent payoff u and that v is at least r. Recasting (2) in terms of $r(\cdot)$ rather than $\mathcal{D}$, the principal's problem is to choose $r(\cdot)$ to maximize

$$\int_0^{u_{\max}} [V(r(u),u)+\alpha u][1-G(r(u),u)]f(u)\phi'(x(u))\ du \tag{3}$$

subject to the "equation of motion"

$$x'(u)=[1-G(r(u),u)]f(u) \tag{4}$$

and the endpoint condition

$$x(u_{\max})=1\ . \tag{5}$$

This optimal control problem is solved formally in the appendix, but its solution can be understood intuitively with the following argument. Consider a point $(u, r(u))$ on the frontier of the permitted set. For the set to be optimal it is necessary that the principal be indifferent between his payoff $[r(u)+\alpha u]$ at that point and his expected payoff from the agent's next-best permitted alternative, conditional on the agent's best permitted project being $(u, r(u))$.

To calculate the latter expected payoff, note that the density that a project drawn at random has payoffs $(u, r(u))$ is $f(u)g(r(u), u)$. Then the probability that one out of n projects has payoffs $(u, r(u))$ and all the others have agent utility no greater than $z \leq u$ or are not permitted is $nf(u)g(r(u), u)[x(z)]^{n-1}$. Taking the sum across n, the probability that one project has payoffs $(u, r(u))$ and all other permitted projects have utility no greater than z is therefore $f(u)g(r(u), u)\phi'(x(z))$. In particular, the probability that the agent's preferred permitted project is $(u, r(u))$ is $f(u)g(r(u), u)\phi'(x(u))$. Conditional on that event, the probability that the next-best permitted alternative for the agent has agent utility no greater than $z \leq u$ is

$$\frac{f(u)g(r(u),u)\phi'(x(z))}{f(u)g(r(u),u)\phi'(x(u))}=\frac{\phi'(x(z))}{\phi'(x(u))}\ ,$$

which has associated density $\phi''(x(z))x'(z)/\phi'(x(u))$. Therefore the indifference (or first-order) condition required for optimality is

$$r(u)+\alpha u=\frac{1}{\phi'(x(u))}\int_0^u [V(r(z),z)+\alpha z]\phi''(x(z))x'(z)\ dz \tag{6}$$

for all $u \in [0, u_{\max}]$.

In particular, we see from (6) that $r(0)=0$. This implies that the principal does not wish to restrict the desirable projects available to the agent whose best project has only zero payoff for her, i.e., there is "no distortion at the bottom". The reason for this is that when $u = 0$ there is no strategic benefit to restricting choice. (The strategic effect of raising $r(u)$ above $-\alpha u$ is to increase the probability that the agent will choose a smaller z, and this effect cannot operate when $u = 0$.) Differentiating (6) with respect to u and using (4) implies the Euler equation for the principal's problem, which is expression (7) below.

Proposition 1 *The principal's problem of maximizing (3) subject to (4) and (5) over all piecewise-continuous threshold functions $r(\cdot)$ has a solution. This solution is differentiable and satisfies the Euler equation*

$$r'(u) + \alpha = [V(r(u), u) - r(u)][1 - G(r(u), u)]f(u)\frac{\phi''(x(u))}{\phi'(x(u))} \tag{7}$$

with initial condition $r(0) = 0$. A sufficient condition for a threshold function which satisfies the Euler equation to be a global optimum is that

$$\zeta(x) \equiv \frac{\phi''(x)}{\phi'(x)} \text{ weakly decreases with } x\ . \tag{8}$$

Proof. The proofs of this and subsequent results are given in the appendix. ■

Expression (7) reveals that ζ in (8) is important for the form of the solution. A short list of examples for this term includes:

- If n is known to be exactly $N \geq 1$ for sure (so $q_N = 1$), then $\phi(x) = x^N$ and $\zeta(x) = (N-1)/x$.

- If n is Poisson with mean μ (so $q_n = e^{-\mu}\frac{\mu^n}{n!}$ for $n \geq 0$) then $\phi(x) = e^{-\mu(1-x)}$ and $\zeta(x) \equiv \mu$.

- If n is Binomial (the sum of N Bernoulli variables with success probability a) then $\phi(x) = (1 - a(1-x))^N$ and $\zeta(x) = a(N-1)/(1 - a(1-x))$. The "known n" case is a special case of the Binomial with $a = 1$. The Poisson is a limit case of the Binomial when $aN = \mu$ and $a \to 0$.

- If n is Geometric (so $q_n = (1-a)a^{n-1}$ for $n \geq 1$ and some parameter $a \in (0,1)$) then $\phi(x) = (1-a)x/(1-ax)$ and $\zeta(x) = 2a/(1-ax)$.

Assumption (8), which states that $\phi'(x)$ is a log-concave function, is valid for the Binomial distribution—and hence for the "known n" and Poisson sub-cases—but not for the Geometric distribution.

Define the "naive" threshold rule to be $r_{naive}(u) = -\alpha u$. This is the threshold rule which permits all desirable projects, i.e., those projects such that $v + \alpha u \geq 0$. This rule might be implemented by a principal who ignored the strategic effect that the agent will only choose the project with the highest u whenever she has a choice. As such, the naive rule is optimal for the principal when the agent never has a choice of project, i.e., when $q_0 + q_1 = 1$. (In this case $\phi'' \equiv 0$, the right-hand side of (7) vanishes, and so $r(\cdot) \equiv r_{naive}(\cdot)$ is optimal.)

Outside this dull case, though, the right-hand side of (7) is strictly positive. Therefore, since $r'(u) + \alpha > 0$ and $r(0) = 0$ it follows that $r(u) > r_{naive}(u)$ when $u > 0$. Therefore, the principal forbids some strictly desirable projects (and never permits an undesirable project). Moreover, the gap between the optimal and the naive rule, $r(u) - r_{naive}(u)$, strictly increases. We state this formally as:

Corollary 1 *Suppose the agent sometimes has a choice of project (i.e., $q_0 + q_1 < 1$). Then it is optimal for the principal to forbid some strictly desirable projects, and the gap between the optimal threshold rule $r(u)$ and the naive threshold rule $r_{naive}(u)$ widens with u. In particular, when $\alpha = 0$ the optimal threshold rule increases with u.*

What is the intuition for why the principal wishes to exclude some desirable projects from the permitted set, whenever the agent sometimes has a choice of project? Suppose the principal initially allows all desirable projects, so that $r(u) \equiv r_{naive}(u)$. If the principal increases $r(\cdot)$ slightly at some $u > 0$, the direct cost is approximately zero, since the principal then excludes projects about which he is almost indifferent (since $r(u) + \alpha u = 0$). But there is a strictly beneficial strategic effect: there is some chance that the agent's highest-u project is excluded by the modified permitted set, in which case there is a chance that she chooses another project which is permitted, say with $z < u$. This alternative project is unlikely to be marginal for the principal, and the principal will expect to get payoff $V(r(z), z) + \alpha z$, which is strictly positive when $r(z) = -\alpha z$. This argument indicates that it is beneficial to restrict desirable projects, and not to permit any undesirable projects. Moreover, it is intuitive that the strategic effect is more important for higher u, since it applies over a wider range $z < u$, and this explains why $r(u) - r_{naive}(u)$ increases with u.

We next discuss some comparative statics for this problem. First, we show that a greater fraction of projects should be permitted when more weight is placed on the agent's interests.

Proposition 2 *Let α_L and α_H be two weights placed by the principal on the agent's payoff, where $\alpha_L < \alpha_H$. Let $r_i(\cdot)$ and $x_i(\cdot)$ solve the Euler equation (7) when $\alpha = \alpha_i$ for $i = L, H$. If assumption (8) holds then $x_L(0) \geq x_H(0)$, i.e., the fraction of permitted projects increases with α.*

Thus we see that the more the principal cares about the utility of the agent, the more discretion—in the sense of a greater fraction of projects being permitted—the agent is given.[11] This result is similar to the "ally principle" in the Holmstrom-type models mentioned in section 1, where the more likely the agent's preferences were to be close to the principal's, the more discretion the agent was given. In those situations in which the agent's salary can be traded off against uncertain on-the-job benefits, the parameter α reflects the terms of that trade-off (which in turn may reflect the degree of the agent's risk-aversion). Proposition 2 shows that when on-the-job benefits are better able to replace salary, the more discretion should be offered to the agent.

A second way in which the ally principal might be expected to be seen concerns the extent of correlation between u and v. Intuitively, when u is positively correlated with v, the agent's incentives are likely to be aligned with those of the principal. In the limit of perfect positive correlation, since the agent's best project is always the principal's best project, it is optimal to give the agent complete freedom to choose a project. (By contrast, with strong negative correlation, the agent's best permitted project is likely to be the principal's worst permitted project.) However, it is not obvious how formally to define a notion of "more correlation" which could be used as basis for general comparative statics analysis. Instead, in section 3.2 we discuss an example which confirms this intuition.

It is also intuitive that when the agent is likely to have more projects to choose from, the principal will further constrain the permitted set of projects. With more

[11] It is not necessarily the case that the threshold rules are *nested* so that $r_H(\cdot) \leq r_L(\cdot)$, and one can find examples where the two threshold rules cross for some positive u.

projects available, the agent is likely to have at least one which lies close to the principal's preferred project. A notion of "more projects" which ensures that this intuition is valid is the familiar *monotone likelihood ratio property* (MLRP). The details are provided in the next result:

Proposition 3 *Let $(q_0^L, q_1^L, q_2^L, ..)$ and $(q_0^H, q_1^H, q_2^H, ..)$ describe two probability distributions for the number of projects which satisfy the monopoly likelihood ratio property, i.e., that q_n^H/q_n^L weakly increases with n. Let $\phi_i(\cdot)$ be the PGF corresponding to $(q_0^i, q_1^i, q_2^i, ..)$, and suppose $\phi_L(x)$ satisfies (8). Let $r_i(\cdot)$ and $x_i(\cdot)$ solve the Euler equation (7) when the number of projects is governed by $(q_0^i, q_1^i, q_2^i, ..)$ for $i = L, H$. Then $x_H(0) \geq x_L(0)$.*

Thus, the fraction of permitted projects falls when the agent is likely to have more projects available. The requirement that the number of projects is ordered by MLRP is a stronger requirement than first-order stochastic dominance. Indeed, there are examples where stochastic dominance leads to a *smaller* fraction of projects being excluded.[12] Moreover, it is not necessarily the case that the principal benefits when the agent has access to more projects. When there is strong negative correlation between u and v, an agent choosing from more projects is likely, all else equal, to choose a worse project from the principal's perspective.

Without making further assumptions, it is hard to make more progress in characterizing the solution to (7). In the remainder of section 3 we examine the properties of the solution in three special cases.

3.1 Independent payoffs and $\alpha = 0$

Suppose that the distribution of v is independent of u and that $\alpha = 0$. Then the principal does not care about the agent's choice of u, either directly (since $\alpha = 0$) or

[12] An example where adding more projects widens the optimal set of permitted projects is the following. Suppose initially the agent has no projects at all with probability $1 - \varepsilon$ and exactly two projects with probability ε. Because the state when no projects materialize plays no role in the determination of $r(\cdot)$, the optimal threshold rule for this agent is just as if there were two projects for sure. Such a threshold rule will strictly exclude some desirable projects. Consider next the situation in which the agent has exactly one more project than the previous situation (i.e., $n = 1$ with probability $1 - \varepsilon$ and $n = 3$ with probability ε). Whenever ε is small, the state where there is only one project will dominate the choice of $r(\cdot)$, and almost all desirable projects will be permitted, thus widening the set of permitted projects. One can check that this pair of probability distributions does not satisfy MLRP.

indirectly since the distribution of his payoff v does not depend on u.

Write $G(v)$ and $V(r)$ as functions which do not depend on u. It follows that

$$\begin{aligned}\frac{d}{du}\left[\frac{V(r(u))-r(u)}{f(u)}\frac{d}{du}\phi(x(u))\right] &= \frac{d}{du}\left[(V(r(u))-r(u))(1-G(r(u)))\phi'(x(u))\right] \\ &= \frac{d}{du}\left[\left(\int_{r(u)}^{v_{\max}}(1-G(v))dv\right)\phi'(x(u))\right]\end{aligned}$$

$$= -r'(u)(1-G(r(u)))\phi'(x(u)) + (V(r(u))-r(u))(1-G(r(u)))^2 f(u)\phi''(x(u))$$

which equals zero at the optimum from (7). Therefore, the second-order Euler equation reduces in this case to the first-order equation

$$[V(r(u))-r(u)]\frac{d}{du}\phi(x(u)) = kf(u) \tag{9}$$

for some positive constant k.

Here, the principal obtains the same expected payoff with all density functions $f(\cdot)$ for u. To see this, change variables in (9) from u to $F(u)$. That is to say, write $\hat{r}(F(u)) \equiv r(u)$ and $\hat{x}(F(u)) \equiv x(u)$, so that $\hat{r}$ represents the threshold rule expressed in terms of the cumulative *fraction* of u-projects F. Then (9) becomes

$$[V(\hat{r}(F))-\hat{r}(F)]\frac{d}{dF}\phi(\hat{x}(F)) \equiv k\ , \tag{10}$$

with endpoint conditions $\hat{r}(0)=0$ and $\hat{x}(1)=1$. Here, the optimal threshold rule $\hat{r}(\cdot)$ does not depend on the distribution for u, as long as u is continuously distributed.[13] As such, only *ordinal* rankings of u matter for the principal in this case.

3.2 Exponential distribution for v

Suppose next that v given u is exponentially distributed on $[0,\infty)$ with mean $\lambda(u)$, so that $G(v,u) = 1-e^{-v/\lambda(u)}$. Suppose that $\alpha = 0$. Since $V(r,u) \equiv r + \lambda(u)$ in this example, the Euler equation (7) is

$$r'(u) = \lambda(u)\frac{d}{du}\log\phi'(x(u))$$

with initial condition $r(0)=0$. Since we wish to compare policies across different distributions for (u,v), the threshold rule r is not in itself insightful. Rather, we study

[13] Note that this argument requires us to change variables in expression (9), and so $F(u)$ needs to be differentiable and, in particular, the distribution for u has no "atoms". If there were atoms, then we would need to consider what project the agent would choose in the event of a tie, when there were two projects which yielded the same maximal agent payoff u.

the *fraction* of permitted projects, and given r write $p(u) = 1-G(r(u), u) = e^{-r(u)/\lambda(u)}$ for this fraction. Writing the Euler equation in terms of p rather than r implies that

$$\frac{d}{du}\log p(u) = -\left[\frac{d}{du}\log\phi'(x(u)) + \frac{\lambda'(u)}{\lambda(u)}\log p(u)\right] \tag{11}$$

with initial condition $p(0) = 1$.

Consider first the case where λ is constant, so that u and v are independent. Expression (11) implies that $p\phi'(x)$ does not vary with u, and it follows that $\phi(x(u)) = k_1F(u)+k_2$ for constants k_1 and k_2. Since $\phi(x(u_{\max})) = 1$, it follows that $k_1+k_2 = 1$. Since $p(0) = 1$, it follows that $k_1 = \phi'(x_0)$, where $x_0 = x(0)$ is the fraction of banned projects at the optimum when u and v are independent. In sum, at the optimum $x(\cdot)$ satisfies $\phi(x(u)) = 1 - \phi'(x_0)(1 - F(u))$. Evaluating this at $u = 0$ implies that the fraction of banned projects satisfies

$$\phi(x_0) + \phi'(x_0) = 1\ . \tag{12}$$

Next, suppose that u and v are positively correlated in the (strong) sense that λ increases with u. Write $h(u) \equiv p(u)\phi'(x(u))$, which from (11) is an increasing function. It follows that

$$\phi(x(u)) = 1-\int_u^{u_{\max}} h(\tilde{u})f(\tilde{u})\,d\tilde{u} = 1-h(0)(1-F(u))-\int_u^{u_{\max}}[h(\tilde{u})-h(0)]f(\tilde{u})\,d\tilde{u}\ .$$

Since $h(0) = \phi'(\tilde{x}_0)$, where $\tilde{x}_0$ denotes the fraction of banned projects in this case with positive correlation, and h is increasing, it follows that $\phi(\tilde{x}_0)+\phi'(\tilde{x}_0) < 1$. Since $\phi(\cdot)+\phi'(\cdot)$ is an increasing function, it follows that the fraction of permitted projects is higher with positive correlation than with independence. A parallel argument establishes that when there is negative correlation, in the sense that λ decreases with u, the fraction of permitted projects is smaller than with independence. In this exponential example, then, positive correlation between u and v is associated with a greater number of permitted projects than negative correlation.[14]

3.3 Poisson distribution for the number of projects

As our third special case suppose that the number of projects follows a Poisson distribution with mean μ, in which case the Euler equation (7) reduces to a first-

[14] One could refine this argument to show that "more correlation" is associated with more projects being permitted. Here, the appropriate notion of "more correlation" is that λ'/λ is shifted upwards.

order differential equation in $r(u)$:

$$r'(u) + \alpha = \mu[V(r(u), u) - r(u)][1 - G(r(u), u)]f(u) \ . \tag{13}$$

The comparative statics of $r(\cdot)$ with respect to α and μ are stronger than the corresponding results in the general setting reported above in Propositions 2 and 3:

Proposition 4 *With a Poisson distribution for the number of available projects, the optimal threshold rule $r(\cdot)$ is pointwise decreasing in α and increasing in μ.*

To obtain some explicit solutions for the threshold rule, suppose that (u, v) is uniformly distributed on the rectangle $[0, 1] \times [-1, 1]$. In this case, (13) becomes the homogeneous equation

$$r' = \tfrac{1}{4}\mu(1 - r)^2 - \alpha \tag{14}$$

with initial condition $r(0) = 0$. Note that if $\mu = 4\alpha$ then the solution to (14) is simply the flat rule $r(u) \equiv 0$. Thus, in the merger context, if the regulator wishes to maximize total welfare (so $\alpha = 1$), then if the expected number of mergers is $\mu = 4$ the regulator should optimally enforce a consumer welfare standard.

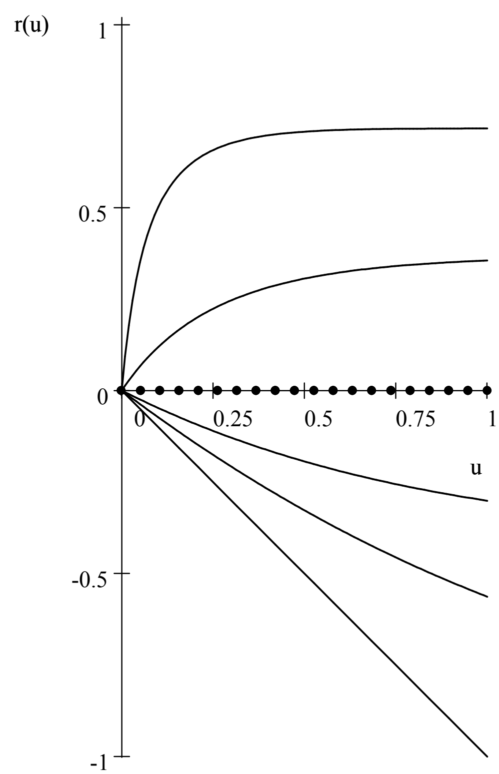

Figure 3: Uniform-Poisson case with $\alpha = 1$ and $\mu = 0$, 1, 2, 4 (dotted), 10 and 50

The solution to (14) when $\mu \neq 4\alpha$ is given implicitly by

$$\int_0^{r(u)} \frac{1}{\frac{1}{4}\mu(1-r)^2 - \alpha} dr = u \ . \tag{15}$$

When $\alpha = 0$, expression (15) yields the simple formula

$$r(u) = \frac{\mu u}{4 + \mu u} \ . \tag{16}$$

When $\alpha > 0$ expression (15) can be integrated using partial fractions to give

$$r(u) = \left(1 - \frac{4\alpha}{\mu}\right) \frac{e^{u\sqrt{\alpha\mu}} - 1}{(1 + \sqrt{4\alpha/\mu})e^{u\sqrt{\alpha\mu}} - (1 - \sqrt{4\alpha/\mu})} \ . \tag{17}$$

Figure 3 plots (17) for $\alpha = 1$ and various μ. Here, higher curves correspond to higher μ, which is consistent with Proposition 4. The straight line depicted for $\mu = 0$ is just the naive rule which permits any desirable project.

A final observation about the Poisson distribution concerns the principal's expected payoff, which from expression (6) evaluated at $u_{\max}$ is equal to $r(u_{\max}) + \alpha u_{\max}$. (Recall that the density of the agent's choice of u is $\frac{d}{du}\phi(x(u))$.) For instance, in the uniform example with $\alpha = 0$ where the threshold rule is given by (16), it follows that the Principal's expected payoff is $r(1) = \mu/(4 + \mu)$.

4 Variants of the Benchmark Model

4.1 Private agent benefits

In some situations it is more natural to suppose that the principal cannot observe the agent's benefit u from implementing a project. In such cases, the principal can permit a project conditional only on the v component. A natural policy in this situation is for the principal to permit those projects which yield payoff $v \geq R$, where R is some threshold or "hurdle rate". In contrast to the benchmark model where u is observable, however, it is not guaranteed that a policy of this threshold form is optimal, and it is sometimes optimal to leave gaps in the permitted set of v-projects.[15] This issue is reminiscent of the question of whether interval delegation is optimal in the Holmstrom-type models discussed in the introduction.

[15] Consider the following example. There are three possible kinds of project, (u_1, v_1), (u_2, v_3) and (u_3, v_2), where $0 < u_1 < u_2 < u_3$ and $0 < v_1 < v_2 < v_3$. It cannot be optimal to ban project (u_1, v_1)

A sufficient condition for a hurdle rate policy to be optimal is that u and v are independently distributed. The reason is similar to the proof of Lemma 1: the density of the agent's preferred permitted project depends only the *probability* that a type-u project is permitted, not on which particular values of v are permitted, and so the principal might as well choose the highest-v projects given the probability of a project being permitted.

If one is willing to restrict attention to hurdle rate policies, either because u and v are independent or for reasons of "simplicity", it is a straightforward matter to derive the optimal hurdle rate. From expression (3), the principal's payoff with hurdle rate R is

$$\int_0^{u_{\max}} [V(R,u)+\alpha u][1-G(R,u)]f(u)\phi'(X(R,u))\ du\ , \tag{18}$$

where

$$X(R,u) \equiv 1 - \int_u^{u_{\max}} [1-G(R,z)]f(z)\ dz$$

is the probability that a given project either has agent payoff z less than u or has v below the threshold R, and the scalar R is chosen to maximize (18).

Some results carry over from the benchmark model where u is observed. Similarly to Proposition 2, for instance, the optimal threshold R decreases with α.[16] If u and v are independent and $\alpha = 0$, the principal's payoff with hurdle rate R in (18) simplifies to

$$V(R)[1-\phi(G(R))]\ . \tag{19}$$

If $\phi_H(\cdot)$ and $\phi_L(\cdot)$ are PGFs corresponding to two distributions for the number of projects, by taking logs of (19), it follows that the optimal hurdle rate is higher with ϕ_H than with ϕ_L if $(1-\phi_H(x))/(1-\phi_L(x))$ increases with x, which in turn is implied by the MLRP condition. Thus, as in the benchmark model (Proposition 3), when the agent has "more projects" in the sense of MLRP she should be offered less discretion over the projects she is permitted to implement.

since the agent only ever chooses that project when it is her sole option. (This is like the "no distortion at the bottom" result in the benchmark model in section 3.) Similarly, there is no point in banning project (u_2, v_3) since that yields the highest-v outcome. Thus the only relevant decision for the principal is whether to ban the intermediate v_2 project. However, it is straightforward to find cases where it is optimal to ban project (u_3, v_2), and so to leave a gap in the set of permitted projects. (For instance, so long as the agent sometimes has a choice of project, this is the case whenever v_3 is sufficiently large.)

[16] Noting that expression (18) can be written as $\int_0^{u_{\max}} [V(R,u)+\alpha u]\frac{\partial}{\partial u}\phi(X(R,u))\ du$, the claim follows by revealed preference.

4.2 Paying for a project

Most of our analysis presumes that monetary incentives to choose a desirable project are not available or desirable, for reasons outside the model. In this section we briefly discuss the principal's optimal policy when he can condition the agent's payment on her performance. We will see that, even within the confines of the model, monetary incentives are not always desirable.

First, suppose that the agent is risk-neutral and is able to bear large losses *ex post* (i.e., she is not liquidity constrained). As in most principal-agent models, the principal here is able to attain his first-best outcome with the use of monetary incentives. The first-best outcome for the principal is obtained when (i) he does not restrict the agent's choice of project, (ii) he pays the agent v when a type-(u, v) project is implemented (and allows the agent to keep her benefit u), and (iii) extracts the agent's entire expected surplus from this scheme in the form of a payment to the principal up front. Such a scheme is akin to "selling the firm" to the agent, and gives the agent ideal incentives to choose the best available project while leaving the agent with zero expected rent.[17]

Outside this extreme case, however, the first-best will not be attainable, and there may again be a role for restricting the agent's discretion. Moreover, the use of monetary incentives will not always be optimal for the principal. To illustrate most simply, consider the situation in which the agent is liquidity constrained in the sense that she must receive a non-negative salary (excluding her payoff u from an implemented project) in all outcomes.[18] For simplicity, suppose that there are just two possible kinds of project, one of which is preferred by the agent while the other is preferred by the principal. Specifically, the "bad project" has payoffs (u_H, v_L) and the "good project" has payoffs (u_L, v_H), where $0 < u_L < u_H$ and $0 < v_L < v_H$. (In this example with perfect negative correlation, the distinction between observable and unobservable u is irrelevant since the principal can infer u from the observed v.) Write $\Delta_u = u_H - u_L$ and $\Delta_v = v_H - v_L$ and suppose $\Delta_v > \Delta_u$ so that (u_L, v_H) is indeed socially the good project. Suppose that any given project is bad with probability P. For simplicity, suppose that $\alpha = 0$ and that there are N projects for sure. If the

[17] This scheme works regardless of whether or not u is observed by the principal.

[18] A similar restriction to non-negative payments is made in Aghion and Tirole (1997), Alonso and Matouschek (2008, section 8.1) and Krishna and Morgan (2008).

principal bans the bad project, his payoff is

$$\pi_1 = (1 - P^N)v_H \ .$$

If the principal allows both projects (without monetary incentives), his payoff is

$$\pi_2 = (1 - P)^N v_H + (1 - (1 - P)^N)v_L \ .$$

(Here, the agent will choose the bad project whenever such a project is present in the N available projects.) The remaining policy is to give the agent a monetary incentive equal to Δ_u to choose the good project (and not to fetter her discretion), which entails payoff

$$\pi_3 = P^N v_L + (1 - P^N)(v_H - \Delta_u) \ .$$

(Here, the agent will implement the good project whenever such a project is available.)

For monetary incentives to be optimal, we require that $\pi_3 \geq \max\{\pi_1, \pi_2\}$. Now

$$\pi_3 - \pi_1 > 0 \Leftrightarrow P^N > \frac{\Delta_u}{v_L + \Delta_u}$$

and

$$\pi_3 - \pi_2 > 0 \Leftrightarrow \frac{(1-P)^N}{1-P^N} < \frac{\Delta_v - \Delta_u}{\Delta_v} \ .$$

These inequalities are jointly satisfied when Δ_u is sufficiently small or when P is sufficiently large. By contrast, monetary incentives should not be used when P is small, Δ_u is large, v_L is small, or N is large.[19]

This example illustrates a more general trade-off between banning mediocre projects and rewarding the choice of good ones. When he bans mediocre projects the principal suffers the cost that such projects are not implemented when they are the only ones available. Rewarding the choice of good projects avoids this cost, but instead involves paying the reward whenever at least one good project is available. Restricting choice is therefore preferred when the chance of having only mediocre projects is small, which is more likely to be true when the agent can choose from many projects.[20] In richer settings than the illustrative binary example above, it may be optimal both

[19]See Figure 9 in Alonso and Matouschek (2008) for an illustration in their framework of the limited gains to the principal in being able to make (non-negative) contingent payments to the agent rather than just banning projects.

[20]Another reason why monetary incentives are not always given to an agent is that the agent performs several tasks, and giving incentives to do one task well might induce the agent to underperform on other, unmeasured, aspects of her job (see Holmstrom and Milgrom (1991)).

to ban mediocre projects and reward the choice of good projects. In addition, if the agent is not liquidity constrained, it is possible to financially penalize her choice of bad projects, which could well be preferable to an outright ban. We leave a more complete analysis of the interactions between restricting choice and monetary incentives for further work.

4.3 Incentives to find a project

The benchmark model in section 3 assumed that the number of projects was exogenous to the agent (but uncertain). In such a framework the agent does not need to be given an incentive to discover projects. In this final variant we suppose that the agent needs to exert effort to find a project. We do this in the simplest possible way, so that by exerting effort e the agent finds a single project with probability e, while with remaining probability $1-e$ no project emerges.[21] If she finds a project, that project's characteristics (u, v) are realized according to the same density functions f and g as in the benchmark model. To achieve success probability e the agent incurs the private cost $c(e)$. Here, $c(\cdot)$ is assumed to be convex, with $c(0) = c'(0) = 0$ and $c'(1) = \infty$.

Since the agent's incentives to find a project depend on her *expected* payoff from a project, her attitude towards risk is relevant (unlike in the benchmark model), and in this section we assume that the agent is risk neutral. (Alternatively, the agent could be risk-averse and u represents her von Neumann Morgenstern utility from the project.) The principal's payoff is a weighted sum of the agent's payoff (including her cost of effort) and the expected value of v, where the relative weight on the agent's payoff by the principal is $\alpha \leq 1$. The principal determines a piecewise-continuous function $r(\cdot)$ such that any (u, v) project with $v \geq r(u)$ is permitted.

If she discovers a project, the agent's expected payoff is

$$A = \int_0^{u_{\max}} u[1 - G(r(u), u)] f(u)\, du\ , \tag{20}$$

[21] A richer model would involve the agent being able to affect the expected number of projects, so that the agent may end up with a choice of project. (For instance, if the number of projects follows a Poisson distribution, the agent could choose μ by incurring cost $C(\mu)$, say.) The principal's optimal policy in this situation has some similarities to the policy when the number of projects was exogenous: the threshold rule is nonlinear and involves $r(0) = 0$. However, like the model of costly discovery analyzed in this section, the threshold rule reflects the need to give the agent an incentive to find more projects (which typically benefits the principal as well as the agent), and it may be optimal *ex ante* to permit projects which are undesirable *ex post.*

and the agent will choose effort e to maximize her payoff $eA - c(e)$. Clearly, a reduction in the threshold rule $r(\cdot)$ induces a higher value of A in (20), which in turn leads to greater effort from the agent.[22] Since high effort benefits the principal as well as the agent, the principal has a reason (beyond the weight α placed on the agent's interests) to increase the leeway given to the agent.

If we write

$$\sigma(A) \equiv \max_{e \geq 0} : eA - c(e)$$

for the agent's maximum payoff given A, then σ is a convex increasing function and $\sigma'(A)$ is the agent's choice of effort e given her reward A. The principal chooses $r(\cdot)$ to maximize his expected payoff

$$\alpha\sigma(A) + B\sigma'(A) \ , \tag{21}$$

where

$$B = \int_0^{u_{\max}} \left(\int_{r(u)}^{v_{\max}} vg(v,u)dv \right) f(u) \ du \ .$$

The principal's optimal policy is described in the next result:

Proposition 5 *The principal's optimal policy takes the form*

$$r(u) + \left[B\frac{\sigma''(A)}{\sigma'(A)} + \alpha \right] u \equiv 0 \ . \tag{22}$$

Proof. Let $r(\cdot)$ be the candidate optimal threshold rule, and consider the impact on the principal's payoff in (21) of a small variation $r(\cdot) + t\eta(\cdot)$ where $\eta(\cdot)$ is an arbitrary piecewise-continuous function. Writing (21) in terms of t yields

$$\Phi(t) = \alpha \int_0^{u_{\max}} u[1 - G(r(u) + t\eta(u), u)]f(u) \ du \ +$$

$$\sigma' \left(\int_0^{u_{\max}} u[1 - G(r(u) + t\eta(u), u)]f(u) \ du \right) \left(\int_0^{u_{\max}} \left(\int_{r(u)+t\eta(u)}^{v_{\max}} vg(v,u)dv \right) f(u) \ du \right) ,$$

and so

$$\Phi'(0) = -(B\sigma''(A) + \alpha) \left(\int_0^{u_{\max}} \eta(u)ug(r(u),u)]f(u) \ du \right)$$
$$-\sigma'(A) \left(\int_0^{u_{\max}} \eta(u)r(u)g(r(u),u)]f(u) \ du \right) \ .$$

Since $\Phi'(0)$ must equal zero for all $\eta(\cdot)$, it follows that $r(\cdot)$ must satisfy (22) ■

[22] This is akin to the "initiative effect" of delegation in Aghion and Tirole (1997).

Thus, the optimal threshold rule is a ray emanating from the origin. Moreover, the ray is downward sloping and weakly steeper than the principal's naive rule (which has slope $-\alpha$). The only situation in which the principal implements his naive rule, i.e., $r(u) = -\alpha u$, is when $\sigma'' = 0$, which applies when the agent's success probability does not respond to incentives, i.e., there is an *exogenous* success probability e. Outside this dull case, though, the principal allows some projects which are strictly undesirable ($v + \alpha u < 0$) in order to stimulate the agent's effort.[23] The more that the agent responds to incentives (in the sense that the function $\sigma''(\cdot)/\sigma'(\cdot)$ is shifted upwards), the more leeway she should have to choose a project. This distortion is the opposite to the bias in the "choosing" model in section 3, where the principal forbade some desirable projects.

Some intuition for the, perhaps surprising, linearity of $r(u)$ comes from the following argument. The principal's payoff in (21) is a function of both A (the expected value of v from a single project given that the project is only implemented if it is permitted) and B (the expected value of u from a project given that the project is only implemented if it is permitted). The problem of maximizing a (nonlinear) function of A and B has the same first-order condition as maximizing a linear sum $A + \gamma B$ for some constant γ. That is to say, the solution to the principal's problem is obtained by choosing $r(\cdot)$ to maximize

$$\int_0^{u_{\max}} \int_{r(u)}^{v_{\max}} [v + \gamma u] g(v, u) f(u) \, dv du$$

for some constant γ, the solution to which is clearly to set $r(u) = -\gamma u$ so that only the positive $[v + \gamma u]$ are contained in the integral.

In earlier work we analyzed a more complicated version of this problem in which the agent searches sequentially for a satisfactory project, and can influence the arrival rate of new projects by incurring effort. In that alternative framework, the agent might not propose the first permitted project which emerges. Rather, she will wait until she finds a permitted project which achieves a reservation utility, where this reservation utility will depend on the threshold rule $r(\cdot)$ as well as her discount

[23] This feature is also seen in Aghion and Tirole (1997) and Baker, Gibbons, and Murphy (1999). By contrast, as discussed in section 1, Szalay (2005) presents a model where information gathering incentives are enhanced by forbidding projects which both the principal and agent might often wish to implement.

rate. In that alternative framework, a linear threshold rule is also optimal for the principal, although not necessarily a rule which starts at the origin. When the principal and agent are more impatient, the threshold rule is shifted downwards, so that the principal is willing to accept a less good project, but with less delay. In the limit of extreme impatience, the dynamic search problem essentially reduces to the framework discussed in this section where the agent tries to discover a single project.

5 Conclusions

Proceeding from the motivating examples of welfare standards in merger policy and project choice within organizations, we have explored the nature of optimal discretion for a principal to give to an agent when the agent may have a choice of project. The principal's problem is to design the optimal set of permitted projects without knowing which projects are available to the agent, though being able to verify the characteristics of the project chosen by the agent. In other words, the problem is to set the optimal rule that the agent must obey, in circumstances where the principal can just check whether or not the rule has been met.

In the benchmark model the agent has a number (unknown to the principal) of projects to choose from. The optimal permission set excludes some projects that are desirable for the principal because the loss from excluding marginally desirable projects is outweighed by the expected gain from thereby inducing the choice of better projects. We showed (i) the principal permits more types of project when he puts more weight on the agent's welfare, and (ii) the principal permits fewer types of project when the agent has more projects to choose from.

In one variant of this model, we supposed that the principal could not observe the agent's payoff from the selected project. Some results from the benchmark model (e.g., concerning comparative statics) were shown to apply in this modified setting, although in contrast to the benchmark model it might be optimal to leave "gaps" in the permitted set of projects. In a second variant, the principal was able to offer a monetary reward to the agent for choosing a good project, but with liquidity constraints on the agent it might nevertheless be preferable to ban mediocre projects than to reward good ones: the former policy has costs when all available projects are mediocre, while the latter involves payments whenever there is at least one good project. In a final variant, by incurring a private cost the agent makes it more likely that a project emerges, and the optimal permission set was characterised by a linear relationship between the payoffs of principal and agent. In order to encourage agent initiative, the principal permits some projects which are undesirable *ex post*, in contrast to the bias induced in the benchmark model.

Our analysis could perhaps be extended to situations where the principal has richer information, e.g., about unchosen projects. For instance, consider modifying the benchmark model so that the principal can cheaply verify the characteristics of *all* announced projects, in which case the agent can only conceal projects but not lie about the projects she has available. In this case, the mechanism design approach no longer coincides with a policy of choosing a permitted set of projects. To illustrate, suppose the principal knows for sure that the agent has $N \geq 2$ projects to choose from. Then, if able to inspect all projects for free, the principal could require the agent to bring all N projects forward, and simply implement his preferred one and so obtain his first-best outcome. Even though such "forcing" mechanisms cannot be used when the number of projects is uncertain (as is more plausible), it would be interesting in future work to analyze the mechanism design approach in this context, and if feasible to derive the principal's optimal rule for selecting one project from the agent's announced list of available projects.

It would also be useful to examine more systematically than we do here the relative benefits of offering financial inducements (including penalties as well as rewards) to choose good projects versus banning mediocre projects. Another way to develop the analysis could be to multi-agent settings: it is after all a feature of many rules that they apply without discrimination to various agents in various situations.

APPENDIX

Proof of Proposition 1:

The principal's aim is to maximize (3) subject to the endpoint condition (5) and equation of motion (4). We consider the control variable $r(\cdot)$ to be taken from the set of piecewise continuous functions defined on $[0, u_{\max}]$ which take values in $[v_{\min}, v_{\max}]$, in which case $x(\cdot)$ is continuous and piecewise differentiable.

Although this is already a well-posed optimal control problem, it is more convenient to consider $s(u) \equiv \phi(x(u))$, rather than $x(u)$, as the state variable. In this case the equation of motion (4) becomes

$$s'(u) = (1 - G(r(u), u))f(u)\tau(s(u)) \ , \tag{23}$$

where $\tau(\cdot)$ is the function derived from $\phi(\cdot)$ such that $\phi'(x) \equiv \tau(\phi(x))$ for all $0 \leq x \leq 1$. (That is to say, $\tau(s) \equiv \phi'(\phi^{-1}(s))$.) Note that τ is an increasing function, and it is (weakly) concave in s if and only if $\phi''(x)/\phi'(x) \equiv \tau'(\phi(x))$ weakly decreases with x.[24] In sum, we wish to maximize

$$\int_0^{u_{\max}} [V(r(u), u) + \alpha u][1 - G(r(u), u)]f(u)\tau(s(u)) \ du \tag{24}$$

subject to the endpoint condition $s(u_{\max}) = 1$ and equation of motion (23). We proceed in three stages: (i) we show that an optimal solution exists; (ii) we derive necessary conditions for the optimal policy, and (iii) subject to a regularity condition we show that any policy which satisfies the necessary conditions is a globally optimal policy.

First, that an solution to problem (24) exists can be deduced from the Filippov-Cesari Theorem (for instance, see Seierstad and Sydsaeter, 1987, chapter 2, Theorem 8). The only non-trivial requirement for this theorem to be invoked is that the set

$$N(s, u) = \{([V(r, u) + \alpha u][1 - G(r, u)]f(u)\tau(s) - \gamma, [1 - G(r, u)]f(u)\tau(s)) \\ : \gamma \geq 0, \ v_{\min} \leq r \leq v_{\max}\}$$

be convex for each s and u. Write $\eta(p, u) \equiv \int_{r(p,u)}^{v_{\max}} vg(v, u)dv$, where $r(p, u)$ is defined implicitly by $G(r(p, u), u) \equiv 1-p$. Thus $r(p, u)$ is the threshold such that a proportion p of projects lie above $r(p, u)$ for given u, and $\eta(p, u)$ is the integral sum of v above this threshold. Therefore, $pV(r(p, u), u) = \eta(p, u)$. Note that η is concave in p, and that the above set N is equal to

$$N(s, u) = \{([\eta(p, u) + \alpha up]f(u)\tau(s) - \gamma, pf(u)\tau(s)) : \gamma \geq 0, \ 0 \leq p \leq 1\}$$

which is therefore convex since $\eta(p, u) + \alpha up$ is a concave function of p. Therefore, an optimal strategy exists.[25]

[24] For instance, in the Poisson case $\tau(s) = \mu s$, and if there are two projects for sure then $\tau(s) = 2\sqrt{s}$. In general, $\tau(1)$ is equal to the expected number of projects.

[25] Strictly speaking, the Filippov-Cesari Theorem shows the existence of a optimal *measurable* control $r(u)$ rather than a piecewise-continuous control. However, in practice this is not an important limitation. (See Seierstad and Sydsaeter, 1987, chapter 2, footnote 9.)

Second, we describe the necessary conditions which must be satisfied by the optimal policy. *Pontryagin's Maximum Principle* (see Seierstad and Sydsaeter, 1987, chapter 2, Theorem 1) states that if a piecewise-continuous control variable $r(\cdot)$ solves problem (24) then there exists a continuous and piecewise-differentiable function $\lambda(\cdot)$ such that $\lambda(0) = 0$ and for all $0 \leq u \leq u_{\max}$:

$$r(u) \text{ maximizes } (V(r,u) + \alpha u - \lambda(u))(1 - G(r,u)) \text{ for } v_{\min} \leq r \leq v_{\max} \tag{25}$$

and except at points where r is discontinuous

$$\lambda'(s) = (V(r(u),u) + \alpha u - \lambda(u))(1 - G(r(u),u))f(u)\tau'(s) \ . \tag{26}$$

Note that (25) implies

$$r(u) + \alpha u - \lambda(u) = 0 \ , \tag{27}$$

and so $\lambda(u)$ represents the gap between the optimal rule $r(u)$ and the naive rule $r_{naive}(u) = -\alpha u$. Since λ is continuous, it follows that r is itself continuous. Moreover, since $r(\cdot)$ is continuous it follows from the Maximum Principle that $\lambda(\cdot)$ is everywhere differentiable, in which case (27) implies that $r(\cdot)$ is itself everywhere differentiable. Since $\lambda(0) = 0$ it follows that $r(0) = 0$. Combining (26) and (27) yields

$$r'(u) + \alpha = (V(r(u),u) - r(u))(1 - G(r(u),u))f(u)\tau'(s(u)) \ ,$$

which is equation (7) in the text.

Finally, we discuss when a policy satisfying these necessary conditions is a global optimum. The *Arrow sufficiency theorem* (see Seierstad and Sydsaeter, 1987, chapter 2, Theorem 5) shows that the necessary conditions pick out a global optimum if

$$[V(r(u),u) - r(u)][1 - G(r(u),u)]f(u)\tau(s)$$

is concave in s for all u. However, since $[V(r(u),u) - r(u)][1 - G(r(u),u)]f(u)$ is positive, the result follows if τ is concave in s. This is so if and only if (8) holds.

Proof of Proposition 2:

Condition (7) implies that at $u = 0$ and any other u such that $r_L(u) = r_H(u)$

$$\frac{r'_L(u) + \alpha_L}{r'_H(u) + \alpha_H} = \frac{\zeta(x_L(u))}{\zeta(x_H(u))} \ . \tag{28}$$

If $x_L(0) < x_H(0)$, then by assumption (8) $\zeta(x_L(0)) \geq \zeta(x_H(0))$, and so (28) implies that $r'_L(0) > r'_H(0)$. In particular, $r_L(u) > r_H(u)$ for small $u > 0$. If $x_L(0) < x_H(0)$ then $r_L(\cdot)$ must cross $r_H(\cdot)$ at some point. (If r_L were uniformly above r_H then clearly the fraction of prohibited projects with α_L would be greater than with α_H.) Let u^* be the first point above zero where the curves cross. In particular, we must have $r'_L(u^*) \leq r'_H(u^*)$. In addition, we must have $x_H(u^*) \geq x_L(u^*)$ since $x_H(0) \geq x_L(0)$ and $r_H(u) \leq r_L(u)$ for $u \leq u^*$. But then (28) implies that

$$1 > \frac{r'_L(u^*) + \alpha_L}{r'_H(u^*) + \alpha_H} = \frac{\zeta(x_L(u^*))}{\zeta(x_H(u^*))} \geq 1 \ ,$$

a contradiction. We deduce that the curves can never cross, and so our initial assumption $x_L(0) < x_H(0)$ cannot hold.

Proof of Proposition 3:

Let $\zeta_i(x) = \phi''_i(x)/\phi'_i(x)$, where ϕ_i is given in the statement of the proposition. We first show that MLRP implies that $\zeta_H(\cdot) \geq \zeta_L(\cdot)$, which in turn follows if we can show that ϕ'_H/ϕ'_L increases with x. However, the derivative of ϕ'_H/ϕ'_L has the sign of

$$\left(\sum_{n=2}^{\infty} n(n-1)q_n^H x^{n-2}\right)\left(\sum_{n=1}^{\infty} nq_n^L x^{n-1}\right) - \left(\sum_{n=2}^{\infty} n(n-1)q_n^L x^{n-2}\right)\left(\sum_{n=1}^{\infty} nq_n^H x^{n-1}\right) \ .$$

We claim that the coefficient on each power x^N in the above is non-negative. Defining $a_k \equiv (k+2)q_{k+2}^H$ and $b_k \equiv (k+2)q_{k+2}^L$, the coefficient on x^N can be written as

$$\sum_{k=0}^{N}(k+1)(a_k b_{N-1-k} - a_{N-1-k}b_k) = (N+1)(a_N b_{-1} - a_{-1}b_N) + \sum_{k=0}^{N-1}(k+1)(a_k b_{N-1-k} - a_{N-1-k}b_k)$$

$$= (N+1)(a_N b_{-1} - a_{-1}b_N) + \sum_{k=0}^{M}[(N-k) - (k+1)](a_{N-1-k}b_k - a_k b_{N-1-k}) \ , \quad (29)$$

where M is the largest integer not greater than $\frac{N-1}{2}$. The final expression pairs together terms in $(N-1-k)$ with terms in k. Since $\frac{a_k}{b_k}$ is increasing in k by MLRP, $\frac{a_{N-1-k}}{b_{N-1-k}} \geq \frac{a_k}{b_k}$ for all $k \leq M \leq \frac{N-1}{2}$. So every term in (29) is positive. It follows that ϕ'_H/ϕ'_L increases with x.

The rest of the proof follows the same lines (with L and H permuted) as that for Proposition 2.

Proof of Proposition 4:

The method is similar to the proof of Proposition 2. Consider first the impact of increasing μ, and let μ_L and $\mu_H > \mu_L$ be two values for μ. Let $r_L(\cdot)$ and $r_H(\cdot)$ be the corresponding optimal threshold rules. From (13) it follows that at $u = 0$ and any other u such that $r_L(u) = r_H(u)$

$$\frac{r'_L(u) + \alpha}{r'_H(u) + \alpha} = \frac{\mu_L}{\mu_H} < 1,$$

so $r'_L(u) < r'_H(u)$ at all such u. So r_H can never cross r_L from above. We deduce that $r_H(u) > r_L(u)$ for all $u > 0$. The argument for the impact of α on $r(\cdot)$ is similar.

References

Aghion, P., and J. Tirole (1997): "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), 1–29.

Alonso, R., and N. Matouschek (2007): "Relational Delegation," *RAND Journal of Economics*, 38(4), 1070–1089.

——— (2008): "Optimal Delegation," *Review of Economic Studies*, 75(1), 259–294.

Amador, M., I. Werning, and G.-M. Angeletos (2006): "Commitment vs. Flexibility," *Econometrica*, 74(2), 365–396.

Baker, G., R. Gibbons, and K. Murphy (1999): "Informal Authority in Organizations," *Journal of Law, Economics and Organization*, 15(1), 56–73.

Dow, J., and C. Raposo (2005): "CEO Compensation, Change, and Corporate Strategy," *Journal of Finance*, 60(6), 2701–2727.

Farrell, J., and M. Katz (2006): "The Economics of Welfare Standards in Antitrust," *Competition Policy International*, 2(2), 3–28.

Fridolfsson, S.-O. (2007): "A Consumer Surplus Defense in Merger Control," in *The Political Economy of Antitrust*, ed. by V. Ghosal, and J. Stennek, pp. 287–302. Elsevier, Amsterdam.

Holmstrom, B. (1984): "On the Theory of Delegation," in *Bayesian Models in Economic Theory*, ed. by M. Boyer, and R. Kihlstrom. Elsevier, Amsterdam.

Holmstrom, B., and P. Milgrom (1991): "Multi-Task Principal-Agent Analysis: Incentive Contracts, Asset Ownership and Job Design," *Journal of Law, Economics and Organization*, 7(1), 24–52.

Kovác, E., and T. Mylovanov (2007): "Stochastic Mechanisms in Settings without Monetary Transfers: The Regular Case," mimeo, University of Bonn.

Krishna, V., and J. Morgan (2008): "Contracting for Information under Imperfect Commitment," *Rand Journal of Economics*, 39(4), 905–925.

Lambert, R. (2001): "Contracting Theory and Accounting," *Journal of Accounting and Economics*, 32(1-3), 3–87.

Lyons, B. (2002): "Could Politicians Be More Right Than Economists? A Theory of Merger Standards," mimeo, University of East Anglia.

Martimort, D., and A. Semenov (2006): "Continuity in Mechanism Design without Transfers," *Economics Letters*, 93(2), 182–189.

Melumad, N., and T. Shibano (1991): "Communication in Settings with no Transfers," *RAND Journal of Economics*, 22(2), 173–198.

Seierstad, A., and K. Sydsæter (1987): *Optimal Control Theory with Economic Applications*. North Holland, Amsterdam.

Szalay, D. (2005): "The Economics of Clear Advice and Extreme Options," *Review of Economic Studies*, 72(4), 1173–1198.