Decision Processes, Agency Problems, and Information: An Economic Analysis of Capital Budgeting Procedures

Anthony M. Marino
John G. Matsusaka
University of Southern California

Corporations use a variety of processes to allocate capital. This article studies the benefits and costs of several common budget procedures from the perspective of a model with agency and information problems. Processes that delegate aspects of the decision to the agent result in too many projects being approved, while processes in which the principal retains the right to reject projects cause the agent to strategically distort his information about project quality. We show how the choice of a decision process depends on these two costs, and specifically on severity of the agency problem, quality of information, and project risk.

Capital budgeting would be easy in a world without agency and information problems. The decision-maker would simply calculate a project's IRR and compare it to the cost of capital. But in the real world, those providing the funds for investment must rely on self-interested agents to identify projects and provide information on expected returns. As a result, the quality of capital allocation depends on how effective the decision process is in attenuating agency problems and bringing forth accurate information. Corporations employ a variety of decision procedures in practice: Some decisions are fully delegated to division and plant managers (typically, expansion of an existing plant); some decisions require approval of headquarters (typically, construction of a new plant); and other decisions require approval conditional on the nature of the proposal, such as when projects requiring more than $1 million go to headquarters while smaller projects can be approved locally.[1]

The main purpose of this article is to shed some light on the tradeoffs between several commonly used budget procedures. To this end, we

---

[1] These examples are from Marshuetz (1985) and Ross (1986).

develop a model of capital budgeting in which a self-interested, informed division manager (agent) identifies projects and makes proposals to a value-maximizing CEO (principal). The CEO must decide what sort of decision-making authority, if any, to yield to the division manager. When the agent derives private benefits from spending, it is easy to see why the principal would want to retain significant decision rights. If informed, the principal can shut down objectionable projects. In the worst case when the principal is completely uninformed and rubber-stamps the proposal, the principal is no worse off than if the decision had been fully delegated. What is harder to understand is why the principal would ever *give up* any rights to reject a project. Put differently, what is the *cost* of retaining a right to intervene? Aghion and Tirole (1997) among others highlight one potential cost: If the agent can be overruled, he might inefficiently reduce his information-collection effort. We focus on a problem that has received less attention: The agent may distort the information transmitted to the principal if the agent fears about being overruled.[2]

To see how the principal's involvement can be costly, consider two simple decision processes. Both begin with the agent identifying a project and making a proposal. We assume that the agent derives private benefits from spending and therefore is more willing to go forward with a project than the principal is. In the first process the principal fully delegates the decision to the agent, while in the second process the principal retains the right to reject the proposal. Under full delegation, the agent proposes his ideal project and the funds are provided as requested. Under the approval process, the agent may behave more strategically. Since the agent has superior information, the principal will attempt to infer something about the project's quality from the agent's spending proposal. If optimal investment is positively correlated with project quality, the principal will view a large spending proposal as indicative of a high quality project. An agent with a low quality project, then, may propose an excessively large budget in order to gain the principal's approval. When the principal cannot separate good proposals from bad ones based on his information, the principal can be worse off with a veto right because projects become inefficiently large.

We develop a simple model to capture this intuition. A key result is that the principal prefers to delegate in situations where an agent with a low quality project would mimic an agent with a high quality project if approval was required (i.e., when approval rights lead to a pooling equilibrium). The principal prefers to retain approval rights when

---

[2] Another distortion can arise at the implementation stage: Zabojnik (2002) explores inefficiencies that may occur when the agent is forced to implement a project that he particularly dislikes.

mimicking does not occur (a separating equilibrium). The value of delegation then depends on whether or not pooling occurs. Several implications follow, among them: (1) Delegation is optimal for projects with low up-side potential ("routine") while approval is better for those with high up-side potential ("innovative"). The reason is that it is less costly for the agent to mimic a project that is just "a little" better than his ideal project than one that is "much" better. (2) Delegation does not necessarily become better than approval as the agent's preferences move into alignment with the principal's preferences. It is possible for a lessening of the agency problem to make pooling more likely because it requires a smaller proposal to successfully pool.

We also explore more elaborate budget processes. One variant is the common threshold approval process: Projects that cost less than a certain amount are delegated while more expensive projects require the principal's approval. In our model, such an arrangement can be superior to unconditional delegation and approval because it reduces the likelihood of pooling by making it optimal for an agent with a low quality project to separate.

Another variant is to set an upper bound on the amount of investment ex ante: The principal announces that proposals in excess of, say, $10 million will not be considered. In our model, this form of capital rationing always reduces spending when the decision is delegated, but it can increase spending when the principal retains approval rights. By restricting the size of high quality projects, a limited budget makes pooling more attractive for an agent with a low quality project. An implication is that ex ante limits are more effective when coupled with delegated decision-making, and can be counterproductive (increase spending) when coupled with an approval process.

Our article is fundamentally about delegation, and is thus part of the nascent literature on the allocation of authority [e.g., Aghion and Tirole (1997), Dessein (2002), and Harris and Raviv (2002), all with roots in Grossman and Hart (1986)]. The literature to this point has focused on comparing two extreme cases, unrestricted decision-making by the principal and unrestricted decision-making by the agent. Such "on–off" assignments of decision authority are useful for studying certain abstract questions about delegation, but actual budgeting practices typically lie somewhere between these extremes. In most corporations, budgeting involves both the principal and the agent, and the process is structured so that both parties have some decision rights: decision authority is fragmented and conditional. Our article expands on previous work by incorporating these forms of partial delegation into the analysis as well. It seems important to consider partial or conditional delegation processes because they are common in practice, and because, as we show, they often lead to better outcomes than unconditional assignments of decision rights.

Thinking about the problem in terms of decision *processes* rather than just decision rights also brings agenda control considerations into the analysis [Romer and Rosenthal (1979)]. The fact is that budgeting usually begins with the agent making a proposal. The ability to move first can allow the agent to influence the outcome even if the agent's formal decision rights are limited, something that is well known by students of budgeting but which has not received much attention in the theoretical literature.[3] In its emphasis on both information and agenda control problems, our article is close in spirit to the pioneering studies of legislative decision-making by Gilligan and Krehbiel (1987, 1989, 1990). Our article also departs significantly from previous work by focusing specifically on capital budgeting. By casting the problem in specific terms, we are able to put more structure on the payoff functions, and this additional structure leads to a richer set of tradeoffs and empirical predictions than have emerged from the previous, more abstract, literature.

Finally, our article can be seen as extending the stream of research pioneered by Harris and Raviv (1996, 1998) showing how some capital budgeting practices can be understood as solutions to agency and information problems. Their papers emphasize costly auditing of information, whereas we assume that information is inherently "soft" and cannot be verified by a third party. The other important difference is that Harris and Raviv (1996, 1998) apply the Revelation Principle to solve their model. In a paper closer in spirit to ours, Harris and Raviv (2002, p. 3) note, "[i]f the Revelation Principle applies, then it is optimal to specify a decision rule that provides incentives for the two managers to communicate their information fully and truthfully ... In this case, one cannot say anything about the level at which the decision is made." Since we are specifically interested in who makes the decisions, we assume [like Dessein (2002) and Harris and Raviv (2002)] that managers cannot commit to mechanisms other than the handful of processes we see in practice.

The article is organized as follows. Section 1 describes the model. Section 2 develops the tradeoff between two simple decision processes, delegation and approval. Section 3 explores threshold approval and capital rationing. Section 4 considers several extensions of the model: An approval process in which the principal can modify the proposal, an informed principal, and an analogous mechanism design problem. Section 5 concludes.

---

[3] Bower (1970, p. 16) notes the importance of moving first in budgeting: "The notion that the decisions of subordinates are crucial to the choices presented to superiors, that indeed these subordinate decisions often may constitute the true shapers and initiators of corporate commitment, once stated is obvious."

1. The Model

The model features a principal who employs an agent to evaluate projects and make proposals. The principal provides the funding for the investment.

1.1 Sequence of actions

There are three periods. In period 0, the principal adopts a decision process. In period 1, the agent possibly receives information about a project's value, and proposes a level of funding. In period 2, the principal can reject the proposal (unless the decision is fully delegated), and if approved, the investment is made and the project pays off. This "bottom up" sequence is a good approximation of actual corporate budget processes, and introduces agenda control considerations into the choice of decision procedure.[4]

1.2 Information

The underlying "quality" of the project is $\theta \in \{H, L\}$ with probabilities π and $1-\pi$, respectively, where $H > L$, and $E[\theta] \equiv M$. If the project is a new plant, we can think of θ as parameterizing the anticipated demand for its product. The agent has private information — with probability p the agent knows the project's quality. Let $S \in \{L, M, H\}$ indicate the agent's information where M indicates no information. The agent's information is "soft" in the sense that it cannot be verified by a third party [which makes our analysis different from Harris and Raviv (1996, 1998)]. At this point, we assume that the principal is uninformed, knowing only the distribution of θ.

Our emphasis on uncertainty about a project's expected cash flows, and our assumption that the agent has an information advantage is consistent with the findings of an extensive survey literature on capital budgeting practices. For example, corporate managers consistently report that "project definition and cash flow estimation" is the most difficult and important stage of the budgeting process rather than financial analysis, project selection, project implementation, and project review. See Bower (1970), Gitman and Forrester (1977), and Scott and Petty (1984).

1.3 Project return

A project's gross return (cash flow) is $\theta f(I)$, where I is the investment or scale of the project and f is increasing and strictly concave with $f(0) = 0$.

---

[4] See Bower (1970), Scott and Petty (1984), and Taggart (1987), for example. "Planning for capital spending is a process which begins with the operating managers of a business. They are the ones who define the needs of their part of the corporation, who make the sales forecasts which justify new capacity, who review technology to determine what the appropriate design should be, who evaluate the economics of a strategy and draft requests for capital funds and, finally, who supervise the design and construction or purchase of a new plant facility and its equipment." [Bower (1970, p. 10)].

The principal provides the funds for the project at a normalized cost of 1 per unit.

1.4 Principal and agent utility functions

The principal and agent are risk neutral. Since the principal receives the cash flow and provides the funds for the project, the principal's utility function is

$$v = \theta f(I) - I. \tag{1}$$

The utility function of the agent is assumed to be

$$u = v + \alpha I, \tag{2}$$

where $0 \leq \alpha < 1$. This formulation has two important features, both of which are fairly standard in the literature. First, the agent cares about the principal's utility, but second, the agent also derives a payoff from project size per se. We shall sometimes refer to α as the severity of the agency problem.

The agent's utility function can be restated as

$$u = \theta f(I) - (1 - \alpha)I. \tag{3}$$

A comparison of Equations (2) and (3) indicates that the principal and agent in our formulation differ only in their private opportunity cost of funds. The agent's opportunity cost of a unit of I is $1 - \alpha$ and the principal's cost is 1. The consequence of this specification is that the agent prefers a larger I than the principal does, other things equal. Note that although the agent wants to over-invest, the agent does not have an unlimited demand for investment. We treat the payoff functions, Equations (2) and (3), as primitives, but think of them as reduced forms arising from a contracting problem that does not perfectly solve the agency problem.[5]

We often calculate the principal's and agent's expected utilities conditional on beliefs about the value of θ. It is convenient to express these expected utilities as v_S and u_S, where S is the expected value of θ conditional on a person's information. For example, u_H is the agent's utility conditional on knowing that the quality of the project is H. When the agent has no information, his utility is u_M.

---

[5] For example, suppose net cash flow is $x = \theta f(I) - I$, the agent is paid a linear contract of the form $a + bx$, and the agent derives private benefits of sI from investment. Then the principal's payoff is $V^1(I,\theta) = x - (a+bx) = (1-b)\theta f - (1-b)I - a$ and the agent's payoff is $U^1(I,\theta) = a + bx + sI = a + b\theta f - (b-s)I$. Since preferences are preserved under affine transformations, we can restate the principal's payoff as $V^2(I,\theta) = \theta f - I$, and the agent's payoff as $U^2(I,\theta) = \theta f - (1 - s/b)I$. We end up with Equations (2) and (3) where $\alpha = s/b$.

The principal and agent disagree about the optimal scale of any project that is approved. To create an interesting conflict, we also assume that they may disagree about whether or not a project is worth funding at all. Specifically, we focus on parameter configurations such that they would both like to go forward (for some I) if the project is known to be high quality, and they both want to shut it down if it is low quality. The disagreement arises when there is no information ($S = M$): The agent would like to proceed but the principal would like to stop. The formal statement of the assumptions is this:

Assumptions: Principal's utility function — $v_L < 0$ for all I, $v_M < 0$ for all I, and $\max_I v_H > 0$; agent's utility function — $u_L < 0$ for all I, $\max_I u_M > 0$, and $\max_I u_H > 0$.[6]

Our use of only three information states is not essential for the results that follow. What is important — and what is different from the more abstract delegation models of Gilligan and Krehbiel (1987), Dessein (2002), and Harris and Raviv (2002) — is the existence of projects of such low quality that the principal is unwilling to fund them at any scale. The existence of this "no compromise" zone emerges naturally when thinking about capital budgeting, and with it, Dessein's result that delegation is always optimal does not appear. We expand on these points below.[7]

2. Two Simple Decision Processes: Delegation and Approval

To highlight the basic tradeoffs, we begin by comparing two simple decision processes. The first is (complete) *delegation*: The agent is given the power to go forward with the project at whatever scale he chooses, and cannot be overruled by the principal. The second is *approval*: The agent proposes a scale and the principal can either approve it without modification or reject it completely (later we will show that nothing of substance changes if the principal can approve the proposal in a modified form). Both processes are common in capital budgeting. For example, Ross (1986) and Taggart (1987) note that decisions about adding capacity for existing products are typically delegated to division and plant managers. Proposals to introduce new products usually require approval at a higher level. Bower (1970, p. 65) emphasizes the up-or-down nature of the approval process: "The (executive committee) review varied in thoroughness depending in large measure on the extent of the project's

---

[6] The assumptions also can be stated: $Lf'(0) < 1 - \alpha < Mf'(0) < 1 < Hf'(0)$.

[7] It may be worth noting that we generate a "no compromise" zone by assuming that the marginal product of the first unit of investment is sufficiently low (see note 6). Alternatively, we could assume that $f'(0) = \infty$ and there is fixed setup cost.

controversialism, but always the result of the review was 'go' or 'no go.' The definition of a project did not change."

2.1 Complete delegation

Under complete delegation, the project goes forward at the agent's optimal scale. Let I_S^* be the optimal investment for the agent [the maximizer of Equation (3)] conditional on the agent's information, S. When nonzero, the optimal investment solves $Sf'(I_S^*) = 1 - \alpha$, and is increasing in S. By assumption, $I_L^* = 0$. The principal's (period 0) expected utility under complete delegation (D) is then

$$E_0[v|D] = (1-p)v_M(I_M^*) + p\pi v_H(I_H^*). \tag{4}$$

Our assumptions in Section 2 imply that the first term is negative and second is positive.

2.2 Approval

Under the approval process, the principal can reject the proposal. The principal will do so if he infers from the proposal that the agent has no information. The uninformed agent ("M-agent") takes this into account when making a proposal. In particular, the M-agent may propose the investment/scale that an H-agent would have chosen, that is, the M-agent may pool with the H-agent.

   A number of different outcomes are possible depending on the parameter configuration, but the interesting economics can be seen by comparing equilibria in which agents pool with those in which they separate. The most transparent cases attain when the principal is willing to accept the H-agent's optimal project size conditional on knowing that the agent has $S \in \{M, H\}$. Therefore, we assume that $v_R(I_H^*) > 0$, where $R = (\pi H + (1-p)(1-\pi)L)/(1-p+p\pi)$ is the expected project quality conditional on $S \neq L$. Given this assumption, there are two Perfect Bayesian equilibria distinguished by one simple condition.

Pooling equilibrium: When $u_M(I_H^*) > 0$, the H-agent and M-agent both propose I_H^*, the principal accepts a proposal of I_H^*, and the principal rejects all other proposals.

   This is an equilibrium because no agent type gains from making a different proposal, and the principal cannot do better with an alternative adoption strategy. The proof is straightforward. Obviously, the H-agent, who is receiving his globally optimal outcome will not deviate. The M-agent's payoff is positive in equilibrium, but zero if he deviates from the equilibrium because his proposal will be rejected. Finally, the principal's behavior is optimal along the equilibrium path because $v_R(I_H^*) > 0$, and his rejection of proposals off the equilibrium path is optimal if he

believes those deviations come from an *M*-agent, which is the only reasonable conjecture.[8]

The important feature of this equilibrium is that an *M*-agent asks for a larger capital allocation than he really wants in order to mislead the principal about the project's prospects. This is somewhat counterintuitive: the principal knows that the agent is excessively fond of spending, but the agent is worried about making *too small* of a proposal. The agent's incentive to "over propose" when the principal is involved in the decision drives the key tradeoffs in the model.[9]

In equilibrium, the project goes forward at a scale of I_H^*, if the agent's information is *M* or *H*. The principal's expected payoff under the approval process (*A*) in this equilibrium is then

$$E_0[v|A] = (1-p)v_M(I_H^*) + p\pi v_H(I_H^*). \tag{5}$$

Separating equilibrium: When $u_M(I_H^*) < 0$, the H-agent proposes I_H^*, the M-agent proposes $I \neq I_H^*$, the principal accepts a proposal of I_H^*, and rejects all other proposals.

The proof is identical to the one above, except that here the *M*-agent would rather not have the project at all than operate it at the *H*-agent's preferred scale. In equilibrium the project goes forward only if the agent knows that $S = H$. Then the principal's expected return is

$$E_0[v|A] = p\pi v_H(I_H^*). \tag{6}$$

To summarize, there are two possible equilibria, and which one attains depends on whether the uninformed agent earns a positive or negative return from mimicking the *H*-agent's proposal.[10]

2.3 Comparison of delegation and approval processes

Now we compare the two decision processes from the principal's point of view. The principal chooses and commits to a decision process in period 0.

---

[8] More formally, this is the only equilibrium that survives the usual refinements. For example, there is another Perfect Bayesian equilibrium where the agents pool at $I_H^* - \delta$, the principal accepts this proposal, and rejects all others. To support this equilibrium, the principal must find it optimal to reject a deviation of I_H^*, which is true only if the deviator's expected type is less than *R* [since $v_R(I_H^*) > 0$]. However, beliefs that give such a large weight to *M* are eliminated by standard refinements such as the intuitive criterion and D1 since this deviation would only benefit the *H*-agent. Equilibria that pool at $I_H^* + \delta$ are eliminated for similar reasons.

[9] An interesting example of proposing overly large projects in order to secure financing comes from the case of now defunct online grocer Webvan. The company opened for business in July 1999, raised $1.2 billion in equity markets, and set out to enter 26 markets before it had figured out how to turn a profit in a single one. It proceeded to lose $100 million a month before liquidating in July 2001. One of the company's VC backers explained, "It's easy to say, 'Man, you could have done a few less markets,' but there was a huge Catch-22. There was a unique opportunity to raise a huge amount of capital in the public market so we could build a business far faster than Sam Walton rolled out Wal-Mart. But to raise money, you had to get above the noise level, build a brand name, and make big promises to investors." Quote taken from, "Some Hard Lessons for Online Grocer," *New York Times*, February 19, 2001.

[10] If $u_M(I_H^*) = 0$, then there are both pooling and separating equilibria.

In practice, it may be difficult for the principal to irrevocably commit to a decision process. We are implicitly assuming that some way to commit is available, such as reputation or repeated play [Baker, Gibbons, and Murphy (1999)].

Casual intuition suggests that the principal would always prefer approval to delegation, since approval entails no opportunity cost. It turns out that delegation is better in some situations.

Proposition 1. *The principal prefers the delegation process when the approval equilibrium pools ($u_M(I_H^*) > 0$), and prefers the approval process when the approval equilibrium separates ($u_M(I_H^*) < 0$).*

The proof follows from comparison of Equations (4), (5), and (6). The intuition is this: Under both delegation and approval, the project goes forward in the H state at scale I_H^* and does not go forward in the L state. The difference appears in the M state. In this state, the principal's payoff is negative for any $I > 0$, and increasingly so as I rises. Under delegation, the project is implemented at a scale of I_M^*. Under approval with pooling, the project also goes ahead, but at an even larger scale, I_H^*, which is worse for the principal. In contrast, under approval with separation, the project does not go forward, which is ideal for the principal.

The basic tradeoff can be summarized as follows: The benefit of approval is that it allows the principal to reject some projects he dislikes; the cost is that the agent will boost proposals to make them appear more valuable. Whether delegation or approval is optimal depends on how willing the agent is to make an exaggerated proposal. Approval is better than delegation when it induces separation and worse when it leads to pooling. The suboptimality of decision processes that induce pooling holds for most of the processes we consider in the article, but not all (see Section 4).

It is natural to wonder whether our results are robust to more complicated information structures than the three-state case we have used here. The results are robust, and the intuitions that emerge from our simple case carry through in a model with an arbitrary finite number of states or with a continuum of states, but the notation is more cumbersome and the intuitions are harder to see. Appendix A explicitly works through the approval equilibrium for both cases: N discrete states and a continuum of states. The reader interested in robustness may wish to skip to that section before proceeding.

In a recent paper, Dessein (2002) also studies the tradeoff between delegation and the approval process (as well as the approval-with-modification process we examine below). In the context of the Crawford–Sobel (1982) model, he proves that delegation is always better than approval (except in the special case where communication is impossible because of extreme agency problems). Our model does not

have this "corner solution" property: Both delegation and non-delegation can be optimal for reasonable parameter values. The source of the difference is not the information structure. As noted, our essential tradeoffs are robust to the cases of N discrete states or a continuum of states. The tradeoff arises in our model because the principal is unwilling to compromise on sufficiently low quality projects, that is, the principal will reject all proposals for very low quality projects. In the Crawford–Sobel model, in contrast, there is always room for compromise.[11] We believe the existence of "no compromise" projects is characteristic of real world budgeting, and our analysis together with Dessein's suggests that the absence of such projects in the Crawford–Sobel model might limit its applicability (at least if one believes, as we do, that "always delegate" is not a good description of the way firms are really organized).[12]

2.4 Implications

The next question is what determines whether delegation or approval is optimal for the principal? Proposition 1 indicates that the answer depends on whether the M-agent pools with the H-agent or separates under the approval process. Formally, delegation is better when $u_M(I_H^*) > 0$. Several observations follow.

1. The approval process becomes better when H rises holding M constant. An increase in H causes I_H^* to rise, which causes $u_M(I_H^*)$ to fall (holding M constant). Intuitively, the increase in H reduces the M-agent utility if he mimics the H-agent's proposal. With a large enough fall in $u_M(I_H^*)$, pooling does not happen, and the approval process becomes optimal. In short, the approval process is more appealing for the principal when the project has a large upside (or variance). One implication is that delegation is better for routine tasks with little upside potential while approval is optimal for new and innovative projects. Bower (1970) notes that returns are easiest to predict for "cost-reducing" projects such as plant modifications and most difficult to predict for projects involving new products. The model implies that decisions concerning plant expansions are more likely to be delegated, while decisions involving new plants and products are likely to be subject to the approval of headquarters—a pattern observed in practice [Ross (1986)].

---

[11] The key is the payoff functions. In the Crawford–Sobel model, the principal and agent have single-peaked utility functions with ideal points of $I = \theta$ and $I = \theta + b$, respectively, where $b > 0$.

[12] We believe that another feature of the Crawford–Sobel model identified by Dessein—optimality of delegation for sufficiently small agency problems—does not appear in our model because of the "no compromise" zone.

2. Casual intuition suggests that as the agency problem becomes more severe, approval is a better choice. This is not necessarily true in our model — an increase in α can make delegation or approval optimal. Intuitively, a rise in α increases I_H^*, which makes pooling less attractive for the M-agent, but it also increases the M-agent's payoffs for a given I. The net effect depends on which of these two forces dominates.[13] The bottom line is that there is not a simple connection between severity of the agency problem and the desirability of delegation.[14]

3. The relation between decision process and project scale is also interesting. Casual intuition suggests that an approval process results in less investment than a fully delegated process. But a simple comparison of the equilibrium outcome reveals that expected investment is higher under delegation than approval, when the approval equilibrium separates (and lower otherwise). Approval can cause spending to go up by inducing the uninformed agent to exaggerate his proposal.

3. Conditional Decision-Making: Thresholds and Capital Rationing

Actual decision processes often employ a mix of delegation and approval depending on the amount of money required. Here we explore two popular examples, threshold approval in which the decision is delegated below a certain amount and requires approval above that amount, and capital rationing in which the decision is delegated below a certain amount and automatically rejected above that amount.

3.1 Threshold Decision-Making

An extremely common practice is to make the decision process conditional on the spending proposal. Most corporations allow division and plant managers to approve small expenditures independently, while a budgeting committee must approve large expenditures [Bower (1970), Ross (1986)]. We call this a *threshold* process and model it as an investment level, T, below which the decision is delegated ($I \leq T$), and above which the project must be approved by the principal.

---

[13] More formally, note that the condition $u_M(I_H^*) > 0$ can be restated as $M/H > \varepsilon(I_H^*)$, where $\varepsilon(I) = If'(I)/f(I)$ is the elasticity of f. An increase in α causes an increase in I_H^*, which can raise or lower ε depending on the precise form of f. One specification in which delegation becomes a better choice when α rises is $f(I) = I - I^2$, with the restriction $I < 0.5$. Here, ε is decreasing in I. If $f(I) = 1 - e^{-1}$ with the restriction $I < 1$, then ε is increasing in I, and an increase in α has the opposite effect.

[14] A similar (ambiguous) result appears in Harris and Raviv (2002) for a Crawford–Sobel (1982) model with quadratic preferences and a uniformly distributed hidden information variable.

We want to identify when a threshold process can be better for the principal than unconditional delegation and approval, and bring out its economic logic. Consider first the situation when the unconditional approval equilibrium pools. The principal faces an unpalatable choice. Under the delegation process, both the M and H projects go forward at the agent's optimum. Under the approval process, both projects will still be approved, but the scale of the M-project will be even larger.

A threshold process can address both problems. To see this, let $\bar{I}$ be the minimum investment that gives the M-agent the same payoff as I_H^*: $u_M(\bar{I}) = u_M(I_H^*)$. The equilibrium with $T \in (\bar{I}, I_H^*)$ is the following: The M-agent proposes $I = \min\{I_M^*, T\}$, the H-agent proposes I_H^*, and the principal approves I_H^* but rejects any other proposal greater than T.[15] The M-agent does not exceed the threshold because any $I \in (\bar{I}, I_H^*)$ results in a higher payoff than I_H^*. The H-agent ends up with his optimal project size and so is willing to accept the principal's oversight instead of proposing a project smaller than the threshold.

How does the principal fare in this situation compared to unconditional delegation and approval? In the H-state, the principal is no better or worse off because the project is funded at I_H^* under each decision process. However, in the M-state the project is smaller than it would be under approval (I_H^*). If the threshold is set below I_M^* then the M-project is smaller than it would be under the delegation process as well. Intuitively, a threshold process addresses the approval "pooling" problem by allowing the M-agent to separate (at a smaller scale) and addresses the delegation "padding" problem by constraining the M-agent (if the threshold is set below I_M^*). A threshold process is obviously worse than an unconditional approval process when the approval equilibrium separates.

Proposition 2. *When the approval equilibrium pools, the principal prefers a threshold process with $T \in (\bar{I}, I_M^*)$ to both unconditional delegation and approval.*

As noted earlier, threshold decision processes are common in practice. One reason is probably because they economize on the principal's time — it is not efficient for the principal to weed out the smallest inefficiencies. Our analysis suggests that a threshold process may have another benefit. By allowing the agent to overspend "a little" on small projects, it prevents even larger distortions that might occur if the agent had to justify his

---

[15] A belief structure that supports this equilibrium is the principal assigning a proposal above T to the H-agent with probability 1. These are the only "reasonable" beliefs when $T \geq I_M^*$. When $T < I_M^*$, a pooling equilibrium in which the threshold is ignored may exist. Since the tradeoffs in that case are already discussed above, throughout this section we assume that the agents play to the separating equilibrium described in the text.

project to the principal. Roughly speaking, a threshold process allows the M-agent to separate while constraining his proposal.[16]

A related question is what determines the optimal threshold? Note that the principal wants to set the threshold as low as possible without inducing the agent to pool, which means the optimal threshold is $T^* = \bar{I}$. Several implications can be derived from the fact that $\bar{I}$ is the solution to $Mf(\bar{I}) - (1-\alpha)\bar{I} = Mf(I_H^*) - (1-\alpha)I_H^*$. First, T^* is decreasing in H, holding constant M. This mirrors our results above: As the project becomes more "routine," the agent is given more discretion. Second, T^* increases as α increases. Somewhat counterintuitively, as the agency problem becomes more severe, it is optimal to give the agent more discretion. The reason is that an increase in α raises the M-agent's payoff from I_H^* more than it raises the payoff from T^*. To prevent pooling, T^* must be increased to make the two payoffs equal again.[17]

3.2 Capital rationing

Another common practice is to limit the total amount of investment ex ante and delegate below that amount, often called "capital rationing" [Gitman and Forester (1977)]. We model this as an upper bound, N, on the available investment. The bound is set in period 0 and cannot be altered thereafter.

Consider a spending limit with delegation first. It is clear that $N > I_H^*$ would have no effect. As N falls below I_H^*, the spending limit cuts the size of the H-agent's project. This makes the principal better off, at least until N reaches the principal's optimal spending level in the H state. Reductions in N below this point will continue to cut investment spending, although this benefits the principal only if the gains from reducing the M-agent's proposal (if any) exceed the losses from reducing the H-agent's proposal.

Now consider an investment limit in the context of the approval process. As above, a limit in excess of I_H^* does not bind. A spending limit below I_H^* reduces the project size in the pooling equilibrium. However, in the separating equilibrium, an investment limit below I_H^* may *increase* the expected project size. This can happen if the limit reduces the H-agent's proposal to the point where the M-agent becomes willing to mimic it, that is, if it transforms a separating equilibrium into a pooling equilibrium. In this case, delegation becomes more desirable than approval for the

---

[16] Harris and Raviv (1998) develop another rationale for something akin to a threshold process based on a model with costly auditing.

[17] One thing missing from our analysis is the possibility that the agent may (inefficiently) subdivide a large project into several smaller projects in order to evade the spending threshold. For example, Bower (1970, pp. 15–16) describes a division that built and equipped an entire plant on expense orders in order to avoid the $50,000 threshold that required approval of top management.

principal. Intuitively, by constraining the *H*-agent, an investment limit makes it harder for the *H*-agent to separate from the *M*-agent. This leads to the next proposition.

Proposition 3. (*a*) *A binding investment limit reduces investment under delegation but can increase investment under approval.* (*b*) *For a sufficiently low investment limit, delegation is always* (*weakly*) *optimal.*

One thing Proposition 3 suggests is that capital rationing and the approval process are substitutes, not complements. In practice, then, we would expect to see capital rationing coupled with delegated decision-making rather than with an approval process. Another empirical implication is that capital rationing is more effective (cuts investment by a larger amount) when used in conjunction with a delegation process than with an approval process.

4. Extensions

We next consider extensions to the model. One purpose is to explore the robustness of the basic tradeoffs.

4.1 The principal can modify the proposal

In the first extension, we allow the principal to modify the agent's proposal instead of only accepting it "as is" or rejecting it outright. This is essentially the process studied in Crawford and Sobel (1982), Aghion and Tirole (1997), Dessein (2002), and Harris and Raviv (2002).

The equilibrium under this type of approval process can display pooling and separation, just as when modification is not possible. To see this, observe that the agent's actual proposal is formally irrelevant so we can think of the agent reporting a state, *L*, *M*, or *H*, and the principal choosing his optimal project size in response. In equilibrium, the *L*-agent reports truthfully and the principal does not proceed with the project. The *H*-agent also reports truthfully; he has nothing to gain by pretending to be an *L*-agent or an *M*-agent. The *M*-agent can either separate (report truthfully) and have the project rejected, or pool (report *H*) and have the project implemented at the principal's optimal scale conditional on $S \in \{M, H\}$, call it I_R^{**}. By definition, $Rf'(I_R^{**}) = 1$. Define I_H^{**} as the solution to $Hf'(I_H^{**}) = 1$.[18] Whether pooling or separating attains equilibrium depends on whether the *M*-agent is willing to mimic: if $u_M(I_H^{**}) > 0$, the *H*-agent and *M*-agent pool and the principal chooses a scale I_R^{**}; if

---

[18] We use one asterisk to indicate the agent's optimal spending levels, and two asterisks to indicate the principal's optimal spending levels.

$u_M(I_H^{**}) < 0$ then the *M*-agent and *H*-agent separate and the principal approves the *H*-project at a scale of I_H^{**}.[19]

The tradeoff between delegation and this type of approval process mirrors Proposition 1, with a few changes in details. As in Proposition 1, the principal prefers approval with changes allowed when the equilibrium separates (in fact, this delivers the principal's first best.) Unlike Proposition 1, however, the principal may prefer approval even when the equilibrium pools. The added benefit comes from cutting back the "padding" that occurs when the agent can make a take-it-or-leave-it proposal. Even so, retaining decision rights is costly for the principal because it causes the agent to distort information, and the principal is better off delegating for some parameter configurations. Note that in all of the comparisons of delegation with variants of the approval decision processes considered to this point (approval, threshold delegation with approval, and capital rationing with approval), the principal prefers the approval decision process to delegation when the *M*- and the *H*-agents separate. Here, we see that this is not the case when the principal can modify the proposal. Whether or not pooling makes for an inferior outcome relative to delegation then is specific to a given process and not a general property.

One thing this clarifies is that delegation does not outperform approval in Proposition 1 because the approval process restricts the principal's ability to react to the proposal. Even if we allow the principal to change the agent's proposal, delegation can still be optimal. The main comparative static implications for the up-or-down approval process also hold for the approval process with changes allowed: Delegation is preferred for projects with low upsides, and the effect of increasing the agency problem on the choice of decision process is ambiguous.

The approval process with changes allowed could be viewed as an alternative to an up-or-down approval process. A natural question is whether one of these processes dominates the other from the principal's viewpoint. The answer, easiest to see by numerical simulation, is no — each can be optimal (and superior to delegation) for some parameter values.[20] Intuitively, the advantage of the approval process with changes allowed is that the principal can cut back the padding by the *H*-agent.

---

[19] The only nonobvious part of the equilibrium is when $u_M(I_R^{**}) > 0$ but $u_M(I_H^{**}) < 0$; both pooling and separating equilibria exist. However, the pooling equilibrium does not survive refinement by the intuitive criterion — when both types propose I_R^{**} in equilibrium and the principal observes an out-of-equilibrium proposal of I_H^{**}, he should believe with probability 1 that the proposal came from an *H*-agent, and accept it.

[20] The least obvious case is where the up-or-down approval process is optimal. Here is a particular example. Let $f(I) = I - I^2$ for $I < 0.5$, and $H = 4$, $L = 0$, $\pi = \alpha = 0.5$, and $p = .8$. This specification meets all of the assumptions and $u_M(I_H^{**}) > 0$ and $v_R(I_H^*) > 0$. The equilibrium under the up-or-down approval process separates, giving the principal an expected payoff of 0.109. The equilibrium under the approval process with changes allowed pools, yielding 0.090. Delegation returns the principal 0.097.

The disadvantage is that pooling is more likely: The *M*-agent is more willing to mimic the *H*-agent when the principal can be relied on to restrict the project's scale.

4.2 Informed principal

So far we have assumed that the principal is completely uninformed about θ. This is a pretty good approximation for many capital budgeting situations. The final decision-maker — the board or an executive committee — has little information about the quality of a project's projected cash flows, cost savings, and so on. Nevertheless, the principal usually has at least a little information and there are cases where the principal might have a great deal of information, such as a proposal to acquire another company.

To get an idea of how sensitive our results are to the assumption of a completely uninformed principal, we worked through an extension of the model in which the principal is informed with probability q. We will not go through the details here because the basic results are easy to describe. Consider the tradeoff between delegation and the approval process with changes allowed. Equilibrium behavior under delegation is the same whether or not the principal is informed. Under the approval process, the *H*-agent and the *L*-agent continue to truthfully reveal their types, and the question boils down to whether the *M*-agent separates or pools with the *H*-agent. When the principal is uninformed, the project will be rejected if the agent reveals that his type is *M*. When the principal is informed, however, the project of the *M*-agent might be approved if the principal's own information reveals that quality is *H*. The upshot is that an *M*-agent is more willing to separate (reveal his type) when the principal is informed than when the principal is uninformed. Otherwise, the analysis of the approval decision process is the same as before.

Two results can be established. First, approval is always optimal for a sufficiently large q. A well-informed principal has little use for the agent's information, and so is willing to risk pooling in order to avoid the padding that occurs under delegation. Second, for sufficiently low q, delegation can be optimal, for the same reasons outlined earlier in the article. In short, we find that the relative information of the principal and the agent affects the decision process in a natural way, and that our main tradeoffs based on information distortion are robust to an informed principal (as long as the principal is not too informed).

4.3 Optimal mechanism from a revelation game

The article focuses on analyzing the benefits and costs of budget procedures that are observed in practice. In this section we investigate how these procedures compare to a theoretically "optimal" decision process. We search for an optimal process using the revelation principle, which allows us to identify optimal mechanisms from among the set of mechanisms in

which the principal is capable of committing costlessly to a specific investment level for each state reported by the agent. This may overstate the mechanisms that are available in practice, since it might be difficult to commit to particularly complicated mechanisms. The exercise is less routine than it first appears in another way: There is no meaningful way to talk about delegation from a mechanism design perspective, since each actor simply reports his information to a machine which then makes a decision [see Harris and Raviv (2002)]. What we are really doing then is finding the optimal mapping between information and investment levels, which we will then compare with the mappings induced by the decision processes studied in the rest of the article.

The revelation principle states that any decision process can be expressed as an equivalent revelation game in which the agent reports a value of θ and is given an incentive to report truthfully. The agent's report, call it $J \in \{L, M, H\}$, results in an investment level. The optimal mechanism is a mapping of reported states into investment that maximizes the principal's expected utility, subject to truth-telling constraints. More formally, it is the I_J defined for $J \in \{L, M, H\}$ that solve:

$$\max_{\{I_J\}} \{p\pi v_H(I_H) + (1-p)v_M(I_M) + p(1-\pi)v_L(I_L)\} \qquad (7)$$

subject to

$$u_J(I_J) \geq u_J(I_K) \quad \text{for all } J, K \neq J, \qquad (8)$$

$$I_J \geq 0 \quad \text{for all } J, \qquad u_J(I_J) \geq 0 \quad \text{for all } J. \qquad (9)$$

Condition (8) imposes truth-telling. Condition (9) contains the non-negativity conditions.

The next proposition (proved in Appendix B) characterizes the solution.

Proposition 4. *An optimal mechanism I_J takes one of three forms depending on the parameters: If $u_M(I_H^{**}) \leq 0$, then (a) $I_L = I_M = 0$ and $I_H = I_H^{**}$; if $u_M(I_H^{**}) > 0$ then either (b) $I_L = 0$ and $I_M = I_H = I_R^{**}$; or (c) $I_L = 0$, $0 \leq I_M < I_M^* < I_H, I_H^{**} < I_H$, and $u_M(I_M) = u_M(I_H)$.*[21]

The optimal mechanism described in Proposition 4, for the most part, can be implemented by the actual decision processes studied in the article. The mechanism in case (a) can be implemented by an approval process with changes allowed. We saw earlier that the approval-with-changes process delivers the principal's unconstrained optimum [of Equation (7)] when the agent separates, which happens when the case (a) holds: $u_M(I_H^{**}) \leq 0$. Case (b) also can be implemented by the approval process

---

[21] A necessary condition for case (b) is $I_R^{**} \leq I_M^*$, but there is not a simple condition to delineate cases (b) and (c).

with changes, although pooling occurs. Here we see an illustration of the point made earlier that a decision process is not necessarily suboptimal just because it induces pooling.

Case (c) is more complicated. The truth-telling condition is difficult to satisfy here, making separation difficult, and the M state is onerous for the principal, making pooling undesirable. The solution is to grant the M-agent a relatively small project, and allow the H-agent a relatively large project. Approval with changes cannot implement such an outcome because the principal is unable to commit to approve such a large project in the H state. Delegation does not work because the M-agent spends too much.

A threshold process (without allowing changes in the proposal) can resolve both of these implementation problems. First, a threshold of $T = I_M$ appropriately caps the M-agent's project size. Second, by granting the agent agenda control power, the principal commits to allow spending in the H state to exceed his personal optimum, I_H^{**}. If $I_H < I_H^*$, a spending limit equal to I_H completes the implementation. If $I_H > I_H^*$, a spending *minimum* is necessary.

A simple approval process (without a threshold) is an optimal mechanism only in the special case where the solution takes the form of case (c) with $I_M = 0$ and $u_M(I_H^*) = 0$.

The only decision process that is never optimal is full delegation. This follows immediately from Proposition 4—the outcomes I_M^* and I_H^* can never occur in an optimal mechanism. However, full delegation is very common in practice. One explanation may be that the analysis omits the opportunity cost of the principal's time. If the principal's time is sufficiently valuable relative to the potential waste from choosing the wrong project size, delegation could be efficient. Still, this argument for delegation seems more applicable to small projects, while large budgeting decisions sometimes are fully delegated as well. Another explanation could involve unmodeled complexity costs. It may be difficult in practice to determine the optimal threshold and spending limits, especially if they vary from project to project and over time, as seems likely.

5. Conclusion

The article studies the economics of several capital budgeting processes that are commonly used by corporations. We develop a model in which the budget process begins with an informed agent making a proposal. The agent prefers to spend more than the principal does, and has superior information about project returns. The principal chooses how much of the decision to delegate to the agent. The central tradeoff is this: Delegation allows the agent to overspend, but when the principal keeps a hand in the decision the agent may distort his proposal to make the project look better

than it is, resulting in an inefficiently large capital allocation. We show how the tradeoff between these two distortions can help explain the choice of decision processes and the behavior of the agent under each process.

One important direction for future research is to investigate the relation of incentive contracts and decision processes.[22] Casual observation and empirical evidence suggest that actual contracts often provide agents with very weak incentives to pursue the principal's interest [Jensen and Murphy (1990)]. It is unclear why this is so. We show that a well-chosen decision process can yield the principal's unconstrained optimal outcome in some cases, so one explanation could be that adroit management of the decision process can address agency problems satisfactorily without having to bear the costs of incentive contracts (such as exposing the agent to significant amounts of risk).

It would also be useful to have a deeper theory of commitment. Our analysis implicitly assumes that the principal can commit to a decision process. Indeed, we argue that some decision processes are effective precisely because they commit the principal to actions that are not in his interest ex post. However, we do not ask why the principal is able to commit to the particular institutions we study and not others. It may well be that some decision processes that are theoretically optimal in a world where commitment is costless (as with a mechanism design framework) are inefficient in reality because of commitment problems.

Finally, the main point of our analysis is that agency and information problems might be useful in understanding how firms choose their budgeting processes. These problems might also be useful in understanding the choice of budget rules [Harris and Raviv (1996)]. It is a longstanding puzzle why so many firms use payback periods and hurdle rates to evaluate projects instead of the theoretically superior net present value technique.[23] We conjecture that one appeal of these popular rules of thumb may be that they are less subject to manipulation by agents, and therefore reduce information corruption.

Appendix A: Generalization to *n* Discrete States and a Continuum of States

This appendix shows that the key features of the approval equilibrium generalize to the case of *n* discrete project types and the case of a continuum of project types. To simplify notation we assume without loss of generality that the agent is always informed, that is, $p = 1$. The uninformed state is unnecessary here because disagreement between the principal and agent can occur in intermediate informed states.

---

[22] Bernardo, Cai, and Luo (2002) makes some progress on this issue.

[23] The popularity of simple rules that do not discount cash flows is well known and enduring. See Graham and Harvey (2001) for recent evidence.
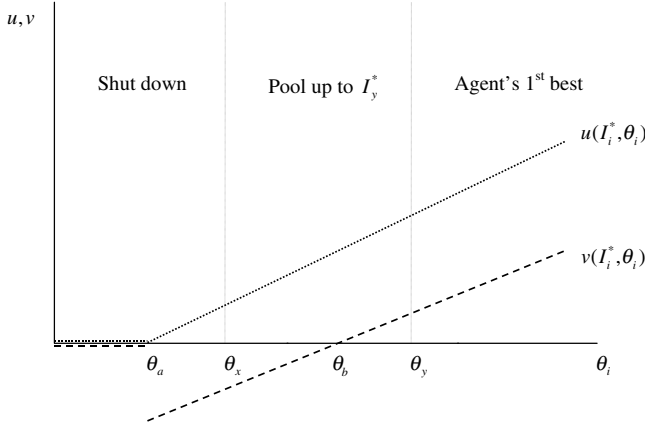
Figure 1
A generalization of approval and delegation.

Consider the discrete case. Let the project types (states) be $\theta_1, \ldots, \theta_n$, ordered so that $\theta_{i+1} > \theta_i$, with probability and distribution functions $g(\theta_i)$ and $G(\theta_i)$, respectively.

The agent's optimal scale in state θ_i, formerly denoted $I_{\theta_i}^*$, is now abbreviated as I_i^*. Recall that $u(I_i^*, \theta_i)$ is nondecreasing in θ_i, and strictly increasing if $I_i^* > 0$. We assume there is a critical value, $a > 1$, such that $u(I_i^*, \theta_i) \leq 0$ for $i < a$, and $u(I_i^*, \theta_i) > 0$ for $i \geq a$.

For the principal, we assume that the following monotonicity condition holds:

$$v(I_i^*, \theta_i) > v(I_{i-1}^*, \theta_{i-1}) \quad \text{for all } i.^{24} \tag{A1}$$

Let $b < n$ be the critical value for the principal such that $v(I_i^*, \theta_i) \leq 0$ if $i < b$, and $v(I_i^*, \theta_i) > 0$ if $i \geq b$. We know that $a \leq b$ because $\alpha > 0$. To create a zone of disagreement between the principal and agent, we assume the inequality is strict: $a < b$.

Given the definitions of a and b, the principal's and agent's payoffs at the agent's optimal scale are as in Figure 1. If the decision is delegated, the agent does not want to go ahead with project types $\theta_i < \theta_a$, but does want to go ahead with project types $\theta_i \geq \theta_a$. Over types $\theta_a, \ldots, \theta_{b-1}$, the principal disagrees with an agent who has decision rights. For types θ_b or greater, the principal is willing to approve the project even at the agent's optimal scale.

To define a perfect Bayesian equilibrium for the approval process, we assume there is a pair (x, y) that satisfies

$$x = \min \{i | u(I_y^*, \theta_i) \geq 0\} \quad \text{and} \quad \sum_{i=x}^{y} v(I_y^*, \theta_i) g(\theta_i) \geq 0.$$

If there is more than one pair, we choose the one with the lowest x. Note that $a \leq x < b \leq y < n$. Let $h(\theta | I)$ be the principal's posterior beliefs conditional on the agent's proposal. The following proposition characterizes a Perfect Bayesian equilibrium of the approval process.

---

[24] In terms of the model parameters, the monotonicity condition boils down to $w(\theta) = -\theta f''/f'f' > \alpha$. One way to satisfy the condition is if w is increasing in θ and there is an appropriate lower bound on θ. Note that w is increasing in θ for a large class of f functions, such as when the degree of concavity, $-f''/f'$, is nondecreasing in I, or when $f''' \leq 0$. So, for example, the monotonicity condition (and our other assumptions) are satisfied if $f = 1 - \exp(-I)$ and $\theta_1 > 1$.

Proposition A1. *Cases* (1)–(3) *below constitute a Perfect Bayesian equilibrium of the approval process that satisfies the Intuitive Criterion.*

1. *Agent proposes $I = 0$ if $i < x$, $I = I_y^*$ if $x \leq i \leq y$, and $I = I_i^*$ if $i > y$.*
2. *Principal approves proposals $I = 0$ and $I = I_i^*$ for $i \geq y$, and rejects all others.*
3. *Beliefs. Along the equilibrium path, $h(\theta_i \mid I = 0) = g(\theta_i)/G(\theta_{x-1})$ for $i < x$, $h(\theta_i \mid I_y^*) = g(\theta_i)/(G(\theta_y) - G(\theta_{x-1}))$ for $i \in \{x, \ldots, y\}$, and $h(\theta_i \mid I_i^* > I_y) = 1$. Off the equilibrium path, $h(\theta_i \mid I_i^*) = 1$ for $i \in \{a, \ldots, b-1\}$, $h(\theta_i = \theta_{b-1} \mid I_i^*) = 1$ for $i \in \{b, \ldots, y-1\}$, and $h(\theta_i = \theta_I \mid I') = 1$ where $I' \neq I_i^*$ for any i.*

Proof. Straightforward comparisons show that the agent and the principal are pursuing Nash strategies given the principal's beliefs. Further, given a proposal off the equilibrium path, the principal's beliefs put zero weight on types that could not benefit from the proposal no matter what the principal does. ∎

The equilibrium features a low quality region where the agent makes no proposal ($\theta < \theta_x$), an intermediate region where the agent pools "up" to I_y^* ($\theta_x \leq \theta \leq \theta_y$), and a high quality region where the agent attains a first best scale ($\theta > \theta_y$). We can understand the tradeoff between delegation and approval by focusing on two particular sets of θ. For $\theta_x \leq \theta \leq \theta_y$, approval induces the agent to pool up. In this region, the principal would be better off delegating in order to avoid the exaggerated proposals. For $\theta_a \leq \theta < \theta_x$, the agent shuts down the project because the agent finds it too costly to pool up, that is, the agent "separates" down. In this region, the principal is better off with approval than delegation because the project does not go forward under approval and it would under delegation. Whether delegation or approval is optimal depends on which of these two effects — "pooling up" or "separating down" — has a larger impact on the principal's utility (in the other regions, the outcome is the same under delegation and approval). The magnitude of the impacts depends on the probability distribution and the parameters of the payoff functions. It can be seen that the tradeoffs here are essentially the same as in our three-state model: Whether approval or delegation is optimal depends on the likelihood of pooling versus separation. This intuition also explains why our model does not produce the "delegation is always optimal" result of the Crawford–Sobel (1982) model that Dessein (2002) highlighted. In our model, the approval process ends up shutting down the low quality projects because there is no project scale that the principal will approve in this region. In the Crawford–Sobel model, there is room to compromise for any quality level, so the agent can always make some proposal acceptable to the principal. As a result, the approval equilibrium ends up with "pooling up" regions but no "separating down" regions, and delegation is always optimal.

A very similar tradeoff appears if we consider the "approval-with-modification" procedure instead of the straight approval process. We will not explicitly work through the details because they are rather involved and we think the intuition extends naturally from Figure 1. Under approval-with-modification, the equilibrium features a shut down region for sufficiently low quality projects, and a series of pooling intervals for higher quality projects. In the shut down region are some projects that would be implemented under full delegation, the analog to $\theta_a \leq \theta < \theta_x$ in the straight approval case; this is where the benefit from *not* delegating appears. Dessein (2002) has shown that the principal would benefit from delegation in the pooling intervals. Approval-with-modification can be optimal if the cost savings from the (marginally) shut down projects exceeds the costs due to scale distortions in the pooling intervals, which can happen for certain parameters. Here again, Dessein's "delegation is always optimal" does not appear because the existence of the no compromise zone for low quality projects.

The extension to the case of a continuum of states is essentially the same. Let $g(\theta)$ and $G(\theta)$ denote continuous density and distribution functions with support $[\theta_L, \theta_H]$, and let $I^*(\theta)$ be

the agent's optimal project size given θ. As before, we assume there is a critical state θ_a such that $u(I^*(\theta), \theta) \leq 0$ if and only if $\theta < \theta_a$. We also retain the monotonicity assumption, $dv(I^*(\theta), \theta)/d\theta > 0$, and assume there is a critical state θ_b such that $v(I^*(\theta), \theta) \leq 0$ if and only if $\theta < \theta_b$. The configuration is then the same as in Figure 1, leading to the same conflicts. A pooling region $[\theta_x, \theta_y]$ can be defined according to Equation (A1) after replacing the summation signs with integrals. Equilibrium of the approval process then takes the same form as proposition A1 with integrals again replacing summation signs.

Appendix B

Proof of Proposition 4. For reference, we restate the mechanism design problem. The optimal mechanism is the I_J defined for $J \in \{L, M, H\}$ that maximizes

$$E[v] = p\pi v_H(I_H) + (1-p)v_M(I_M) + (1-p)(1-\pi)v_L(I_L), \qquad (B1)$$

subject to

$$u_J(I_J) \geq u_J(I_K) \quad \text{for all } J, K \neq J, \qquad (B2)_{JK}$$

$$I_J \geq 0 \quad \text{for all } J, \qquad (B3)$$

$$u_J(I_J) \geq 0 \quad \text{for all } J. \qquad (B4)$$

The optimal mechanism takes one of three forms. We consider them in order.

Solution (a). If $u_M(I_H^{**}) \leq 0$, then the solution is the unconstrained maximum of (B1) and, as we show in the text, it can be implemented (by approval with changes allowed).

Solutions (b) and (c). If $u_M(I_H^{**}) > 0$, then the solution can take two forms, pooling and separating. We begin by establishing two properties of an optimal mechanism. ∎

Lemma 1. *An optimal mechanism satisfies $I_L = 0$.*

Proof. Assume to the contrary that $I_L > 0$. We will show that $I_L' = 0$ yields a higher payoff to the principal and still satisfies the constraints. On the first point, note that $E[v|I_L] < E[v|I_L' = 0]$ because $v_L(I)$ is decreasing for $I \geq 0$. As for the constraints, (B4) is satisfied by $u_L(0) = 0$. Constraints $(B2)_{LK}$ hold by $u_L(0) > u_L(I_K)$ for $K \neq L$. Finally, constraints $(B2)_{KL}$ hold because $u_K(I_K) \geq u_K(0) = 0$ for $K \neq L$. ∎

Lemma 2. *An optimal mechanism satisfies $I_M \leq I_H$.*

Proof. This result follows from the agent's truth-telling constraints. Define $\phi(S) = u_S(I_H) - u_S(I_M)$. The truth-telling constraints can be restated for the H-agent $(B2)_{HM}$ as $\phi(H) \geq 0$, and for the M-agent $(B2)_{MH}$ as $\phi(M) \leq 0$. Now suppose that $I_M > I_H$. Observe that $d\phi/dS = f(I_H) - f(I_M) < 0$. If $\phi(H) \geq 0$, then $\phi(M) > 0$: The two constraints cannot be satisfied simultaneously. ∎

Lemmas 1 and 2 imply that there are at most two types of solutions to the revelation game when $u_M(I_H^{**}) > 0$, those that pool, and those that separate with $I_M < I_H$. We next characterize them and show that a locally optimal solution of each type exists for all parameter values. The proof is completed with a numerical example showing that both pooling and separating solutions can be globally optimal for some parameter configurations.

Solution (b): Pooling. Consider a mechanism with $I_L = 0$ and $I_M = I_H = I'$. For $I_M = I_H = I'$, the payoff function in (B1) can be simplified as $E[v] = (1 - p(1-\pi))Rv_R(I')$. The

unconstrained maximum is $I' = I_R^{**}$. Since $I_R^{**} \leq I_M^*$ by assumption, we have a strictly concave function defined on a compact convex set, $\{(I_M, I_H)|I_M = I_H \leq I_M^*\}$. The maximizer I_R^* exists and is unique.

Solution (c). Separating with $0 \leq I_M < I_M^* < I_H, I_H^{**} < I_H,$ *and* $u_M(I_M) = u_H(I_H)$. We shall first characterize the separating solution(s), supposing they exist. Note that $I_M < I_H^*$. This must be true for (B2)$_{HM}$ to hold. With this in mind, we can show that (B2)$_{MH}$ holds with equality. Again let $\bar{I}$ be the solution to $Mf(\bar{I}) - (1-\alpha)\bar{I} = Mf(I_H^*) - (1-\alpha)I_H^*$. Suppose there is a solution for which this constraint does not bind, then a smaller value of I_M still satisfies (B2)$_{MH}$, and also satisfies (B2)$_{HM}$ because $I_M < I_H^*$. However, the smaller value of I_M increases $E[v]$, which cannot be true for a solution. If $I_M = 0$, a decrease in I_H has the same effect, because $I_H > I_H^* > I_H^{**}$. Next, we observe that $I_M < I_M^* < I_H$. This follows from the fact that (B2)$_{MH}$ binds, $I_M < I_H$, and u is single-peaked. The last property is $I_H > I_H^*$. Suppose that there is a solution for which this is not the case. Then an increase in I_H increases $E[v]$. Constraint (B2)$_{MH}$ holds because $I_H > I_M^*$, and (B2)$_{HM}$ holds because $I_H^{**} < I_H^*$.

Next, we show that a solution of this form exists. Note that the problem can be restated as max$E[v]$ subject to $u_M(I_M) = u_M(I_H)$ for $I_M \in [0, I_M^*)$ and $I_H > I_M$. The set of (I_M, I_H) defined by the constraint is compact if we add the point (I_M^*, I_M^*). Because the objective function is continuous, a maximum exists. The solution is obviously not at the "compactification" point because $I_M < I_M^*$. Although the objective function is strictly concave, the solution is not in general unique. A sufficient condition for uniqueness is that the constraint function $I_H = \psi(I_M)$ defined by $u_M(I_M) = u_M(I_H)$ for $I_M \in (0, I_M^*)$, $I_H > I_M$, is regular convex.[25]

Finally, we show that both pooling and separating solutions can be globally optimal for some parameter values using a numerical example. Let $f(I) = I - I^2$ for $I < 0.5$. Let $H = 4$, $L = 0$, and $\pi = 0.25$, so that $M = 1$. Also, let $\alpha = 0.5$ and $p = .8$. In this case, the pooling solution to (B5) is $I_H = I_M = 0.3$, with $E[v] = 0.09$, whereas the separating solution is $I_H = 0.4$ and $I_M = 0.1$, with $E[v] = 0.11$. Here, the separating solution is the globally optimal mechanism. If α is changed to 0.85, the pooling solution is the same, but the separating solution is $I_H = 0.47$ and $I_M = 0.38$, with $E[v] = 0.07$. Now the pooling solution is the globally optimal mechanism. For each example, it is easy to check that all of the assumptions of our model are met [in particular, $u_M(I_H^{**}) > 0$]. In addition, for this f, the constraint function ψ is convex, so the separating solution we study is the unique maximum. ∎

References

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

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

Bernardo, A., H. Cai, and J. Luo, 2002, "Capital Budgeting and Compensation with Asymmetric Information and Moral Hazard," *Journal of Financial Economics*, 61, 311–344.

Bower, J. 1970, *Managing the Resource Allocation Process: A Study of Corporate Planning and Investment*, Harvard University, Division of Research, Graduate School of Business Administration, Boston, MA.

Crawford, V., and J. Sobel, 1982, "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.

Dessein, W., 2002, "Authority and Communication in Organizations," *Review of Economic Studies*, 69, 811–838.

---

[25] ψ is regular convex if $u''_M(I_M)/u''_M(I_H) \geq (u'_M(I_M)/u'_M(I_H))^2$, or equivalently, $\psi'' \geq 0$.

Gilligan, T., and K. Krehbiel, 1987, "Collective Decisionmaking and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures," *Journal of Law, Economics, and Organization*, 3, 287–335.

Gilligan, T., and K. Krehbiel, 1989, "Asymmetric Information and Legislative Rules with Heterogeneous Committees," *American Journal of Political Science*, 33, 459–490.

Gilligan, T., and K. Krehbiel, 1990, "Organization of Informative Committees by a Rational Legislature," *American Journal of Political Science*, 34, 531–564.

Gitman, L., and J. Forrester Jr., 1977, "A Survey of Capital Budgeting Techniques Used by Major U.S. Firms," *Financial Management*, 6, 66–71.

Graham, J., and C. Harvey, 2001, "The Theory and Practice of Corporate Finance: Evidence from the Field," *Journal of Financial Economics*, 60, 187–243.

Grossman, S., and O. Hart, 1986, "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 691–719.

Harris, M., and A. Raviv, 1996, "The Capital Budgeting Process, Incentives and Information," *Journal of Finance*, 51, 1139–1174.

Harris, M., and A. Raviv, 1998, "Capital Budgeting and Delegation," *Journal of Financial Economics*, 50, 259–289.

Harris, M., and A. Raviv, 2002, "Allocation of Decision-Making Authority," working paper, Graduate School of Business, University of Chicago.

Jensen, M., and K. J. Murphy, 1990, "Performance Pay and Top-Management Incentives," *Journal of Political Economy*, 98, 225–264.

Marshuetz, R., 1985, "How American Can Allocates Capital," *Harvard Business Review*, 63, 82–91.

Romer, T., and H. Rosenthal, 1979, "Bureaucrats Versus Voters: On the Political Economy of Resource Allocation by Direct Democracy," *Quarterly Journal of Economics*, 93, 563–587.

Ross, M., 1986, "Capital Budgeting Practices of Twelve Large Manufacturers," *Financial Management*, 15, 15–22.

Scott, D., Jr., and J. Petty II, 1984, "Capital Budgeting Practices in Large American Firms: A Retrospective Analysis and Synthesis," *Financial Review*, 19, 111–123.

Taggart, R., 1987, "Allocating Capital Among a Firm's Divisions: Hurdle Rates vs. Budgets," *Journal of Financial Research*, 10, 177–190.

Zabojnik, J., 2002, "Centralized and Decentralized Decision-Making in Organizations," *Journal of Labor Economics*, 20, 1–22.