Organizational Design: Decision Rights and Incentive Contracts

Susan Athey and John Roberts*

Where should decision rights be lodged in organizations? Michael C. Jensen and William H. Meckling (1992) argue that moving a decision away from the inherently best-informed party involves costs in communication and garbling but may lodge it with someone who has better incentives to make good decisions. But generally we expect that incentives are part of the organizational design. Why not just provide incentives to those with the best information so that they make the right decisions?

One reason is that the available incentive instruments must serve multiple purposes and designing them to induce better decisions worsens performance against other organizational objectives. Our experience suggests this is a common situation in actual organizations: The means available to affect one sort of behavior or decision inevitably affect the incentives governing other choices. Then the design of incentive schemes and the allocation of decision rights become interlinked.

This paper looks at this idea in the specific context of a principal's problem of inducing agents to provide unobservable effort while also motivating the efficient selection of investments. Each of these problems has been extensively studied in isolation (on inducing effort, see, e.g., Bengt Holmström 1979, Holmström and Paul Milgrom 1991; on decisions see, e.g., Eugene F. Fama and Jensen 1983, Milgrom and John Roberts 1990a, 1990b, Philippe Aghion and Jean Tirole 1997, Tirole and Matthias Dewatripont 1999). We thus know that motivating effort is done best by rewarding agents on precise measures of their effort, not necessarily on the total value created in the firm. At the same time, it is clear that getting the right investment choices may require that the decision makers' rewards be tied to total value created. The difficulty is that the available measures do not allow doing both. The only available performance measures are aggregates whose component pieces cannot be disentangled, while contracts must be written in advance of learning about investment possibilities. Including many contingencies in the contracts *ex ante* is impossible, and on-going renegotiation in every *ex post* eventuality is prohibitively costly.

More formally, we assume that it is not possible to contract on investment projects, nor can the principal bargain with the agents over the adoption of these projects once they are identified. Instead, returns to projects are reflected in the performance measures available for use in the effort-incentive contracting. Then the incentives for effort and for decisions are inextricably tied together. In this framework, we explore the interactions among the design of jobs and assignment of individuals to tasks, the shape and intensity of effort incentives, and the allocation of authority over project selection.

We argue that it may indeed be optimal to assign decisions rights to someone other than the best-informed party. An authority-based hierarchy then emerges endogenously, with some agents being given the right to make organizational decisions over projects that others discovered. Moreover, as in Herbert Simon (1952), those in authority will make the decisions in a self-interested way. Simon emphasized that this will not be *ex post* efficient. By designing the incentives appropriately, however, the inefficiency may be mitigated.

1. The Model

Consider a situation with a risk-neutral principal who has two tasks to be performed, as in Holmström and Milgrom (1991). In addition, investment opportunities (or “projects”) arise. Projects may or may not be undertaken. If a project is not undertaken, the principal’s expected returns (before any payments she makes) are $\Pi_1(e_1) + \Pi_2(e_2)$, where e_i is the total effort devoted to task i and the Π_i functions are concave. If a project is undertaken, the returns are $\Pi_1(e_1) + \Pi_2(e_2) + y_1 + y_2$, where y_1+y_2 is the total return to adopted projects. There are two performance measures, $x_1 = \Pi_1(e_1) + y_1 + \varepsilon_0 + \varepsilon_1$ and $x_2 = \Pi_2(e_2) + y_2 + \varepsilon_0 + \varepsilon_2$, where the ε_i are noise terms, with $v_i = \mathrm{var}(\varepsilon_i)$ and $v_1 \leq v_2$. Assume that the project returns are random, with the two dimensions being identically and independently distributed with $\mathrm{E}(y_i) = 0$, and let the distributions be symmetric about 0. Let ξ_i denote $\mathrm{E}(\ y_i \mid y_i \geq 0)$: This is a measure of the importance of project choice.

There are two agents, A and B, with mean-variance utility functions $\mathrm{U}^i = \mathrm{E}(w^i) - \mathrm{c}(e^i) - r\mathrm{var}(w^i)$, $i = A, B$, where w^i is the agent's wealth, the cost-of-effort function (c) is strictly convex, and r parameterizes the agent's risk aversion. Each agent is willing to accept a contract if his expected utility exceeds his reservation utility. Each agent can allocate effort to either or both tasks: $e_j = e_j^A + e_j^B$, $j = 1, 2$, and $e^i = e_1^{\ i} + e_2^{\ i}$, $i = A, B$.

The signals x_i are observable and contractible, as is commonly assumed in the agency literature, and we assume that payments to the agents are linear functions of the signals.[1] We think of the projects as being numerous, but for simplicity, we will focus on just one project. We assume that exactly one of the two agents "discovers" this project (and each agent is equally likely to do so). When an agent "discovers" a project, he learns the corresponding payoffs y_1 and y_2. These are not contractible, nor is whether a project has or has not been implemented. Note that adopted projects shift the means of the two signals. Thus the dependence of the payments to the agents on the signals determines the agents' preferences over projects. (From an *ex ante* perspective, the decision rules about projects affect the variance of payoffs as well.)

In the spirit of the separation of ownership and control, we assume that the principal cannot learn the returns to projects at any cost. Then the principal cannot participate in the investment decision. It is worth noting, however, that the principal does not automatically have the right incentives for project selection in any case, because she will be concerned not with the total returns but with returns net of payments to the agents.

We also assume that projects cannot be suppressed – if an agent discovers a project then the other agent can, at cost k, observe its returns. This seems to be the appropriate modeling for projects like choosing an advertising agency, a supplier, or a quality level, where some decision must be made. For other sorts, it might be more appropriate to assume that an agent can hide projects he does not like, and the monitoring applies only to projects that the agent reveals. Moreover, we ignore the problem of inducing the agents to bear the cost k to learn the value of the other agent's projects, simply assuming that the agents incur k when the regime calls for it. The analysis is similar but more complicated when these assumptions are relaxed, as discussed in Susan Athey and Roberts (2001).

The principal designs the contracts to maximize her *ex ante* expected profits, which (due to the assumption that utility is transferable through the wage contracts) yields the same solution as maximizing total value.

If there were no investment choice issue, the solution to the effort elicitation problem would be straightforward, at least for v_0 not too small. Assign one agent to each task (even though multitasking is possible) – say, A to task 1 and B to 2. To improve the accuracy of the signals and thus permit more intense incentives and increased effort, agent A's incentive contract places a negative weight on x_2 to reduce the effect of the common noise ε_0, and likewise for agent B. Thus, reward A with $w^A = \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2$ and B with $w^B = \beta_0 + \beta_1 x_1 + \beta_2 x_2$, where $\alpha_2 = -\alpha_1 v_0/(v_0+v_2)$ and $\beta_1 = -\beta_2 v_0/(v_0+v_1)$. As v_0 approaches 0, α_2 and β_1 also go to zero, so that each agent is rewarded solely on the direct signal of his own task; but, even then, it will typically be optimal for the agents to specialize on only one task.[2] Then α_1 and β_2 are chosen subject to the usual incentive constraint that the return to effort must equal the marginal cost, while α_0 and β_0 are set so that the participation constraints are just met. We refer to incentives such as these, with positive weight on one signal and negative weight on the other, as "comparative performance evaluation."

Now consider the induced preferences over project selection. While total surplus is increased by a project if and only if $y_1 > -y_2$, agent *A* will wish to implement a project if and only $y_1 > -(\alpha_2/\alpha_1)y_2$, while *B* will implement if and only if $y_2 > -(\beta_1/\beta_2)y_1$. Note that this means that each agent has very bad incentives on project selection, because each puts negative weight on one element of the returns. Notice further that if $v_2 > v_1$, then $\alpha_1 > \beta_2$ and $\alpha_2/\alpha_1 = -v_0/(v_0+v_2) > -v_0/(v_0+v_1) = \beta_1/\beta_2$. Thus, *A*'s decision-making incentives are less distorted than *B*'s.

On the other hand, if only project selection were important, then the solution would give each agent a payment of $\gamma(x_1+x_2)$ for arbitrarily small but positive γ, which would cause them to value total wealth creation. Then each can be left to adopt whatever projects he sees and likes.

We consider several alternative allocations of decision rights between the two agents over project selection. In the first regime, "empowerment," each agent can choose whether or not to implement any projects he discovers. In the second regime, "*A* authority," agent *A* learns the payoffs of *B*'s potential projects (at cost *k*) and can overrule *B*'s decisions. Thus projects are adopted if and only if they benefit *A*. The third regime, "*B* authority," is analogous to *A* authority. Finally, we consider collegial or "consensus" decision-making, where both agents learn the payoffs from the others' projects and both must agree for a project to go forward. Note that all but the empowerment regime give authority over an agent's projects to the other agent or the two acting together as a committee.

Two other possible allocations of authority are conceivable but do not seem relevant in the context of our model. One is “prohibition,” under which the agents are not allowed to undertake any projects. While this might be attractive in some circumstances, it is hard to see how it might be enforced under the assumptions we have made. The second is a “bargaining” regime, where the agents can make transfers and decide jointly over project allocation. If the bargaining and transfers could be limited to the investments, allowing this option would be attractive. However, it is hard to see how this restriction could be enforced, and the problems created by allowing side payments between agents, particularly under comparative performance evaluation and multitasking, are well known (Holmström and Milgrom, 1990). We discuss bargaining at greater length in Athey and Roberts (2001).

We now proceed to compare the returns to different authority regimes and reward schemes in different parameter regions for our model. The model is sufficiently complex that we do not attempt to characterize optima except in extreme cases, but instead focus on the nature of the solutions, the trade-offs involved, and comparative statics.

2. Motivating Effort and Decisions

We begin with the case where decisions are not very important (so that providing effort incentives dominates the determination of the reward schemes). Suppose that v_0 is large relative to v_1 and v_2. Then the optimal incentives will have strong comparative performance evaluation, with $\alpha_2 \cong -\alpha_1 < 0$ and $\beta_1 \cong -\beta_2 < 0$, and the best allocation of decision rights will be consensus, at least for k not too large. This is because allowing new projects to go forward under the empowerment or the two agent-authority regimes has a mean effect that is close to zero in terms of overall profits, but it leads the agents to bear some extra risk. Since the agents are risk averse, *ex ante* total value is decreased by allowing project adoptions, and the ideal would be prohibition. However, since the agents will be in almost complete disagreement about projects, consensus essentially eliminates all projects and so largely replicates the effect of prohibition (although the cost k is incurred for each project). However, the few projects that are implemented are profitable, so if k is small then consensus will certainly be best.

As v_0 falls for fixed v_1 and v_2, the use of comparative performance evaluation decreases, and so the distortion of the agents' decision-making decreases. At some point, with $v_2 > v_1$, if k is large relative to the agents' risk aversion, a regime of A authority may be preferred. To see this, recall that A's decisions are less distorted than B's, which means that A authority is better than B authority, while consensus involves incurring k on all projects, which is prohibitively costly. As v_0 falls further, however, the empowerment regime may become preferred, because B's decisions also become less distorted, and empowerment saves the cost k on B's projects.

On the other hand, if k is small relative to the agents' risk aversion, consensus decision-making may be preferred even as v_0 approaches zero. Since the use of comparative performance evaluation diminishes as v_0 falls, the acceptance region for agent A approaches $y_1 \geq 0$, while that for B approaches $y_2 \geq 0$. The mean decision quality is the same under consensus decision-making and empowerment, but consensus leads to fewer projects and thus reduced variance in agent payoffs. For small k, the extra cost of k on each project is not sufficient to offset this gain.

Thus, if the importance of the projects is not great, the solution depends on the magnitude of the common noise term and the costs of reviewing project returns. Empowerment, consensus or A authority may be best, depending on the parameters.

Next, we turn to analyze what happens as decisions grow more important, so that it is optimal to modify the reward scheme to account for its effect on decision quality. Starting from the benchmark where v_0 is large and k is small, recall that consensus decision-making may be preferred. In that case, as decisions grow more important the optimal scheme will moderate the commission rates to improve decision quality. This will be implemented by making α_2 and β_1 less negative. In turn, this leads to weaker effort incentives (lower α_1 and β_2), because the variance of payments increases.

Finally, we consider an alternative organizational design that may fare well if decision-making is very important and e_2 is not too important: agent B does not receive any comparative performance evaluation, but instead has $\beta_1 = \beta_2$, and B is given authority over all project decisions. Because agent B now must be exposed to the risk in ε_0 if effort is to be induced, it will be best to set low levels of β_1 and β_2, and not much effort will be induced on task 2. As well, if e_1^A is large enough that $\beta_1\Pi_1'(e_1^A) < \beta_2\Pi_2'(e_2^B)$, then e_1^B will be zero. Thus B is apparently receiving effort incentives for task 1 that are quite costly and yet do not "work," in that no effort is induced. Since all of A's decisions will be reviewed by B, who is a perfect decision-maker, there is no reason to moderate the use of comparative performance evaluation for A: $\alpha_2 = -\alpha_1 v_0/(v_0+v_2)$. Thus, allocating authority to agent B allows the firm to provide agent A with pointed, narrow incentives that are intense but poorly aligned with overall firm value. Agent B's incentives are moderated and his effort is reduced, but his low-powered incentives are aligned with overall value creation, making him a better decision-maker.

This organizational design achieves perfect decision-making and second-best effort on task 1 at the cost of reduced effort on task 2 and inefficient risk bearing by B. Each of agent A's projects must be reviewed at cost k, while agent B's projects don't need to be reviewed. In general, however, it is not optimal to set $\beta_1 = \beta_2$. So long as effort on task 2 is somewhat important, we will have $\beta_1 < \beta_2$, reflecting the fact that β_1 is used only to induce balanced decisions but incurs risk, while β_2 is needed to induce effort on task 2. Thus decision-making is biased: y_2 is overvalued relative to y_1. This organizational design will be optimal if decision-making is important enough, effort on task 2 is not too important, and k is small enough.

3. Hiring a Third Agent

This logic suggests that in general, if k is small enough, it can be useful to give all authority to an agent with balanced incentives. Thus, if the cost of hiring an agent is not too large, when decision-making is very important it may be optimal to use a third agent, C, identical to the other agents, to make decisions. The following organizational design may then be optimal. Assign agents A and B to tasks 1 and 2, respectively, and give them the optimal incentives from the perspective of effort allocation (that is, comparative performance evaluation). Give authority to agent C and pay him $\gamma(x_1+x_2)$, where $\gamma>0$. In this case, agent C makes perfect decisions, while agents A and B receive the effort incentives that are optimal in the absence of decision-making considerations. Thus, the organization is characterized by a "CEO" who receives a small share of overall performance and "division managers" who receive comparative performance evaluation. The CEO specializes in project evaluation and investment decisions, while the division managers exert high levels of effort on managing their respective divisions.

Finally, it is natural to consider the possibility that such a third agent finds it more expensive to monitor project returns than the agents who exert effort in tasks 1 and 2: Agent C incurs a cost $k'>k$ to monitor projects. Consider a regime where $\alpha_1>0$, $\beta_2>0$, $|\alpha_2|< \alpha_1$ and $|\beta_1| < \beta_2$, and decision rights are allocated as follows: Agents A and B can implement, without review, any projects that both agree on, and they can also individually call projects to the attention of agent C. In other words, C reviews all projects where A and B disagree.

Under these assumptions, if k and k' are small enough, conflicting incentives have a desirable effect: Any value-enhancing project that A opposes will be taken to C by agent B, and vice versa. Thus, all good projects come to the attention of agent *C*, who correctly evaluates them. This will prevent good projects from being missed. A potential cost of unbalanced incentives for *A* and *B*, however, is that too many bad projects might be brought forward for C's review. However, anticipating rejection by agent *C*, *A* will not see any gain from appealing *B*'s rejection of projects that increase *A*'s wage but decrease total value. Then it is weakly optimal for *A* and *B* to protest exactly the subset of value-maximizing projects that help one agent and hurt the other. Finally, observe that since *A* and *B* agree only if the projects increase total value, allowing *A* and *B* to implement projects when they have consensus economizes on *C*'s review costs.

Summarizing, if it is more costly for agent *C* to evaluate decisions than *A* and *B*, it may be profitable for *A* and *B* to review one another's projects, and for agent *C* to review only projects where the agents disagree. A scheme with these main features has been used, for example, at IBM Corporation, where an operating unit had to get the concurrence of other such units to implement projects. If concurrence was withheld, the decision was passed to a common hierarchic superior (Richard F. Vance, Arvind Bhambri and James Wilson, 1980).

4. Extensions and Conclusions

Our model has endogenized the simultaneous design of incentive schemes and allocation of decision rights, generating (in some cases) an authority-based hierarchy. Major gaps remain, however: to characterize fully the optima, to allow the agents to choose whether to incur the costs of evaluating others' projects, to allow agents to suppress projects they discover whose adoption would harm them, and to introduce a cost of discovering projects. These are addressed in Athey and Roberts (2001).

References

Aghion, Philippe and Tirole, Jean. "Formal and Real Authority in Organizations." *Journal of Political Economy*, February 1997, *105*(1), pp. 1-29.

Athey, Susan and Roberts, John. "Authority and Incentives in Organizations." Mimeo, Stanford GSB, 2001.

Dewatripont, Matthias and Tirole, Jean. "Advocates," *Journal of Political Economy*, February 1999, *107*(1): 1-39.

Fama, Eugene F. and Jensen, Michael C. "Separation of Ownership and Control." *Journal of Law and Economics*, June 1983, *26*(2), pp. 301-25.

Hart, Oliver and Moore, John. "On the Design of Hierarchies: Coordination versus Specialization," NBER Working Paper No. 7388, 1999.

Holmström, Bengt. "Moral Hazard and Observability," *The Bell Journal of Economics*, Spring 1979, *10*(1), pp. 74-91.

Holmstrom, Bengt and Milgrom, Paul. "Regulating Trade among Agents." *Journal of Institutional and Theoretical Economics*, March 1990, *146*(1), pp. 85-105.

_______, "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics and Organization*, Special Issue 1991, *7*(0), pp. 24-52.

Jensen, Michael C. and Meckling, William H. "Specific and General Knowledge, and Organizational Structure," in L. Werin and H. Wijkander, eds., *Contract Economics*. Oxford: Blackwell Publishers, 1992, pp. 251-74.

MacDonald, Glenn and Marx, Leslie. "Adverse Specialization," *Journal of Political Economy*, forthcoming.

Milgrom, Paul and Roberts, John. "Bargaining Costs, Influence Costs, and the Organization of Economic Activity," in J. E. Alt and K. A. Shepsle, eds., *Perspectives on Positive Political Economy*. Cambridge: Cambridge University Press, 1990a, pp. 57-89.

_____, "The Efficiency of Equity in Organizational Decision Processes," *American Economic Review*, May 1990b (*Papers and Proceedings*), *80*(2), pp. 154-159.

Simon, Herbert. "A Formal Theory of the Employment Relationship," *Econometrica*, July 1951, *19*(3), pp. 293-305.

Vance, Richard F., Bhambri, Arvind and Wilson, James. "IBM Corp.: The Bubble-Memory Incident" Harvard Business School case 9-180-042, 1980.

Notes

---

* Athey: Department of Economics, Massachusetts Institute of Technology, E52-252C, 50 Memorial Drive, Cambridge, MA 02142 (e-mail: athey@mit.edu) and, for AY 2000-2001, National Fellow, Hoover Institution; Roberts: Graduate School of Business, Stanford University, Stanford, CA 94305-5015 (e-mail: roberts_john@gsb.stanford.edu). We are grateful to Bentley MacLeod and Michael Riordan for their suggestions.

[1] Following the logic of MacDonald and Marx (forthcoming), it may be profitable to use nonlinear contracts that induce a complementarity between x_1 and x_2. Although many of the specific predictions of the model would change if nonlinearities were allowed, we would still face the same basic tradeoff as in the current modeling.

[2] Specialization eliminates the fixed cost (through the risk premium) of giving agent incentives for a second task using a second signal (Holmström and Milgrom, 1991). For small v_0 there is another possible solution in which B is given a positive β_1 and is induced to supply both sorts of effort. This might occur if v_2 were large, so that it is very expensive to induce a lot of e_2, leaving B with a low marginal cost of effort. In this case, however, we would actually need $\beta_1 > \beta_2$ if the Π_i functions are similar, so that B's investment preferences over-weight the easily measured task. We will largely ignore this possibility in what follows.