

Robust Incentives for Effort

GABRIEL CARROLL

Stanford University

We discuss a new perspective on the classic principal-agent problem, which asks how to optimally incentivize an agent to exert costly effort on some task, by rewarding him based on his performance. We review recent results on a robust version of the problem, in which the principal is uncertain about what actions the agent can and cannot take and evaluates an incentive contract based on the worst case over possible environments; the worst-case-optimal contracts turn out to take a simple form.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics

General Terms: Economics, Theory

Additional Key Words and Phrases: Incentives, Information Acquisition, Linear Contracts, Principal-Agent Problem, Robustness

1. THE PRINCIPAL-AGENT PROBLEM

Imagine a *principal* who wants to get an *agent* to exert effort on some task. Effort is costly, so the agent won't do it without some incentives. Specifically, the agent can be rewarded as a function of his performance. But performance is an imperfect signal of effort. For example, the agent might be a consultant hired to solve a problem, and the principal cannot observe how hard he tries, only whether he succeeds in solving the problem (or, perhaps, how good a solution he provides). The more effort he puts in, the more likely he is to find a good solution; but he might simply get lucky and find a good solution without trying, or conversely, might fail because the problem is really hard even if he puts in considerable effort. The principal's problem is then as follows: how to design the *contract* $w(\cdot)$, mapping performance y into payment $w(y)$, so as to optimally trade off the goals of inducing effort and not paying too much.

This *principal-agent problem* has vast and diverse applications in economics. Principal-agent models have been used not only to describe employment relations, but also regulation of monopolies, bank lending, corporate governance, and incentives for politicians. There are also many computation-related applications that involve incentivizing costly effort, such as motivating investments in network security, peer-to-peer routing, or crowdsourcing.

The standard modeling approach assumes given a complete description of the environment — utility functions for the principal and agent, and knowledge of all the possible effort levels the agent can exert and the (stochastic) mapping from effort to outcomes. Given this, in a fairly general formulation, one can find the optimal contract $w(\cdot)$ by convex programming [Grossman and Hart 1983]. However, if we want to take this approach literally to design incentive schemes, there are two difficulties. One is, of course, that it is not very realistic to assume the principal

Author's address: gdc@stanford.edu

knows a complete description of the agent’s possible actions. The other is that the optimal contract that comes out is often sensitive to the details of the problem — as often happens with Bayesian mechanism design problems — and can have unintuitive properties, such as payment not being monotonic in output. To roughly summarize the state of the literature, the sufficient conditions known to guarantee a well-behaved optimal contract are fairly restrictive, such as a one-dimensional choice of effort and a monotone-likelihood-ratio assumption on the resulting output distributions [Grossman and Hart 1983; Innes 1990].

On the other hand, when explicit performance-based contracts appear in the real world, they tend to be simple. A particularly common form is *linear* contracts, which pay a fixed fraction (say 10%) of whatever output the agent produces [Bhattacharyya and Lafontaine 1995; Chu and Sappington 2007]. This has posed a puzzle for contract theory [Holmström and Milgrom 1987]: how should the model be written so as to accurately account for the prevalence of such contracts? Alternately put, what do real contract writers know that is missing from the model?

2. ROBUST INCENTIVES

It turns out that uncertainty about the environment is exactly the missing ingredient: Linear contracts provide a guarantee on the principal’s net profit without requiring much knowledge of the environment. To see this, consider risk-neutral, quasilinear preferences — so the agent will choose his action to maximize expected payment minus the cost of effort, and the principal is concerned with expected output minus payment. Assume that the agent can never be paid less than 0. Suppose that the principal knows the agent has some action he can take that produces an expected output of 90 (the exact distribution is not important here) at an effort cost of 10. The principal may know of some other actions as well. But she is not sure she knows all the possible actions; the agent may have tricks up his sleeve that she cannot foresee. Now consider a linear contract that pays the agent $1/3$ of whatever output he produces. By taking this known action, the agent earns an expected payoff of $(1/3) \times 90 - 10 = 20$. So even if the agent takes some other, unforeseen action, his expected payoff will be at least 20 (otherwise this action would not be in his interest). Since the principal pays the agent $1/3$ of output and keeps $2/3$ for herself, her payoff is at least twice the agent’s, and thus her expected payoff is at least 40. So we have a guarantee for the principal.

All well and good; but do linear contracts play a special role in this story? They do indeed, as shown in [Carroll 2013b] in a quite general version of the model. Consider the risk-neutral principal and agent, and consider *any* set \mathcal{A}_0 of possible actions by the agent, where an action is described by a (nonnegative) effort cost and a distribution over (nonnegative) output. \mathcal{A}_0 describes the actions that the principal knows the agent can take. Suppose the principal evaluates any contract by its guarantee — the worst-case expected net profit $E[y - w(y)]$, over all possible true sets of actions \mathcal{A} that contain \mathcal{A}_0 as a subset. Then, the optimal guarantee is always attained by a linear contract. Moreover, one can modify the model to make it more realistic and the conclusion continues to hold. For example, the principal might know that it is impossible to produce expected output y at a cost of less than $b(y)$, where b is some given convex function. With this restriction on possible

actions, the optimal guarantee still comes from a linear contract.

3. INFORMATION ACQUISITION

This same methodology can be fruitfully applied to the problem of studying incentives to acquire information. This is an important problem in reality; there are entire industries built around gathering information. Yet most of the relevant literature on incentives here — such as scoring rules and prediction markets — focuses on getting participants to truthfully report what they already know, not on motivating them to learn more. Indeed, with the usual model of a fully-specified environment, finding the optimal contract has proven tractable only in special cases [Zermeño 2011]. But here too, the problem of optimizing guarantees in an uncertain environment allows us to make inroads.

Suppose there is a finite set of states of the world, and a principal who must make a decision — say, which of several products to produce and sell — whose payoff depends on the state. The principal can hire an expert, who can learn about the state by exerting effort. The expert reports what he learns, the principal makes her decision, and at some later date, the state becomes publicly revealed. The principal can promise to pay the expert based on how accurate his report turned out to be; more precisely, payments can be any function of his report and the realized state (but, as before, cannot be less than zero). The principal knows some of the actions the expert can take to seek information — such an action now being represented by a nonnegative cost of effort and a probability distribution over posterior beliefs about the state, such that the posterior must on average equal the prior belief. As before, we assume that there may be other, unforeseen actions, and a contract is evaluated based on its worst-case guarantee on the principal's expected profit.

A linear contract in this situation would have the expert report his posterior belief, then the principal make her optimal decision, and pay some fixed fraction of her payoff to the expert. Actually, a more general version would include an additive term, which may be positive or negative and could depend on the state. For example, if every decision produces a payoff of at least 100 in state ω , the principal might promise $1/3 \times$ (her payoff, minus 100 if ω is realized). These additive terms allow the principal to pay less on average without affecting the incentives to acquire information.

As shown in [Carroll 2013a], the best guarantee in this setting actually comes from a variant of a linear contract, which specifies a subset D' of the principal's decision set D . The expert is asked to report both the best true decision $d \in D$, as well as the best restricted decision $d' \in D'$. The principal follows decision d , but pays the expert a fraction of the payoff that d' *would* have produced in the realized state, plus some state-dependent additive term. Restricting the set of decisions D' helps the principal because it can allow her to make the additive terms lower without violating the nonnegativity constraint.

4. FURTHER COMMENTS

The work summarized here joins a small but growing literature on how simple mechanisms can provide optimal guarantees in uncertain environments. For example, Frankel [2013] considers the problem of delegating a collection of similar decisions,

such as asking a teacher to grade many students, where the teacher’s preferences about the relationship between grade and actual performance are unknown, and shows conditions under which the optimal mechanism is simply a “budget” specifying how many of each grade should be given. Garrett [2013] considers a procurement problem, in which the principal can see how much the agent spent to produce a good, but not how much effort was exerted to bring the costs down, and shows that the robustly optimal mechanism gives the agent the choice between some fixed price (in which case he gets to keep any further savings) and full-cost reimbursement. Chassang [2013] gives a class of principal-agent environments in which linear contracts are robustly optimal, as here, though his objective is the worst-case *ratio* of the principal’s profit to total social welfare (or equivalently, to the profit the principal could attain if she knew the environment).

Because of the pervasiveness and variety of agency models, there are many opportunities to expand on the work described here, and try to discover similarly simple optimal contracts in other kinds of environments with uncertainty about agents’ actions. For example, future work could explore possibilities with multiple agents, multiple principals, hierarchical agency structures, or actions taken sequentially over time.

REFERENCES

- BHATTACHARYYA, S. AND LAFONTAINE, F. 1995. Double-sided moral hazard and the nature of share contracts. *RAND Journal of Economics* 26, 4, 761–781.
- CARROLL, G. 2013a. Robust incentives for information acquisition. Working paper, Stanford University.
- CARROLL, G. 2013b. Robustness and linear contracts. Working paper, Stanford University.
- CHASSANG, S. 2013. Calibrated incentive contracts. *Econometrica* 81, 5, 1935–1971.
- CHU, L. Y. AND SAPPINGTON, D. E. M. 2007. Simple cost-sharing contracts. *American Economic Review* 97, 1, 419–428.
- FRANKEL, A. 2013. Aligned delegation. *American Economic Review*. Forthcoming.
- GARRETT, D. F. 2013. Robustness of simple menus of contracts in cost-based procurement. *Games and Economic Behavior*. Forthcoming.
- GROSSMAN, S. J. AND HART, O. D. 1983. An analysis of the principal-agent problem. *Econometrica* 51, 1, 7–45.
- HOLMSTRÖM, B. AND MILGROM, P. 1987. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica* 55, 2, 303–328.
- INNES, R. D. 1990. Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory* 52, 1, 45–67.
- ZERMEÑO, L. 2011. A principal-expert model and the value of menus. Working paper, Massachusetts Institute of Technology.