Behavioral Mechanism Design: Optimal Crowdsourcing Contracts and Prospect Theory

DAVID EASLEY and ARPITA GHOSH
Cornell University

Incentives are more likely to elicit desired outcomes when they are based on accurate models of agent behavior. A growing literature—from behavioral economics, as well as online user studies—suggests, however, that people do not quite behave like standard economic agents in a variety of environments, both online and offline. What consequences might such differences have for the optimal design of these environments?

In this note, we summarize our results from Easley and Ghosh [2015] which explores this question of behavioral design—how departures from standard economic models of agent behavior affect mechanism design—via the problem of the optimal choice of contract structure in crowdsourcing markets where workers make decisions according to prospect theory preferences [Kahneman and Tversky 1979] rather than classical expected utility theory: we show that a principal might indeed choose a fundamentally different kind of mechanism—an output-contingent contest versus a 'safe', output-independent, fixed-payment scheme—and do better as a result, if he accounts for deviations from the standard economic models of decision-making that are typically used in theoretical design.

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

General Terms: Mechanism design, Behavioral economics, Game theory, Crowdsourcing

1. INTRODUCTION

The vast range of systems with outcomes that depend on the choices made by economic agents has led to a rich and large literature on mechanism design, which regards designing incentives so that agents make choices resulting in 'good' outcomes. Incentives are more likely to elicit desired outcomes when they are derived based on accurate models of agent behavior. A growing literature, however, suggests that people do not always behave like the standard economic agents in the mechanism design literature. Can such differences have significant consequences for the optimal design of incentive mechanisms?

Ghosh and Kleinberg [2014] investigate this question of 'behavioral' mechanism design in the context of the optimal design of contests, a commonly used incentive mechanism in crowdwork, for *simple agents*: unlike agents who strategically choose the quality of their submission to the contest, simple agents only strategize about *whether* or not to participate in a contest, but not over the quality of their submission. But why, or rather when, a contest? In Easley and Ghosh [2015], we explore another dimension of deviation from the standard economic models of behavior via this question: what kind of incentive structure should a principal in a crowdsourcing environment use—a contest, a fixed payment scheme, and so on—and how does the answer to this question depend on how agents make decisions under uncertainty?

Email addresses: `dae3@cornell.edu, arpitaghosh@cornell.edu`

Decision making under uncertainty: Prospect theory. Many environments to which the mechanism design literature has been applied involve risky choice—agents who must make decisions between a set of choices, each of which yields associated payoffs with some probability. The standard economic model for choice under uncertainty is *expected utility theory* [Neumann and Morgenstern 1944; Savage 1954], whereby agents make choices by comparing the expected utility of outcomes resulting from each possible choice. A substantial body of work in behavioral economics, however, demonstrates that individuals display systematic deviations from the standard expected utility model, including overweighting low-probability events and under-weighting high-probability ones, as well as displaying a greater disutility for losses compared to utilities from a gain of the same amount.

Prospect theory, introduced in Kahneman and Tversky [1979] and further refined to cumulative prospect theory in Tversky and Kahneman [1992], is a descriptive model that describes much empirically observed decision-making behavior more accurately than expected utility theory, and can explain a wide range of experimentally documented behavioral biases, including the status quo bias and endowment effects. One of the best-known achievements of behavioral economics, prospect theory led to the award of the 2002 Nobel prize in economics to Kahnemann for his " ... contributions to behavioral economics, in particular the development of cumulative prospect theory".

Yet the literature on mechanism design almost uniformly models agents as making choices according to the tenets of the classical expected utility theory. While the expected utility model may accurately model choice-making for the applications addressed by classical mechanism design—such as auction theory, where large firms, which are possibly exempt from these behavioral biases, are the decision-making agents—a number of newer online systems such as crowdsourcing and labor markets, peer-to-peer economies, and online auctions with small bidders, involve individual agents to whom these behavioral biases do apply. What consequences might such biases in decision-making have for the design of these environments?

2. OPTIMAL CONTRACT STRUCTURE IN CROWDSOURCING MARKETS

In Easley and Ghosh [2015], we explore this question in the context of principal-agent problems in online labor and crowdsourcing markets. Consider a principal with a single task that he wants to complete by hiring workers via an online platform. The principal has many options available in terms of what incentive structures he can use, such as those supported by various online crowdsourcing platforms—fixed-price contracts on platforms such as Amazon Mechanical Turk (MTurk) or their analog with price discovery, an auction to determine the payment for the task as on oDesk, as well as contests of various kinds on platforms like TopCoder or Kaggle. How should the principal choose between these various kinds of incentive structures—and does the answer to his question depend on whether his population of potential workers behaves according to the classical model of expected utility theory, or whether they deviate, behaving instead according to prospect theory?

We consider a simple model with a principal with a single task and a population of potential workers or agents, each of whom incurs an opportunity cost c to undertaking the task. Agents make endogenous participation choices, but quality is exogenous[1] *i.e.*, participation, but not quality, is a strategic choice: each agent who strategically decides to undertake the task produces output with quality $q_i \sim F(q)$ at a cost c. Agents evaluate each possible action choice according to the value of the corresponding *prospect* it induces: a prospect which offers payoffs x_k with probabilities p_k is evaluated according to $\sum_k u(x_k)\pi(p_k)$, where u is the utility function for payoffs and π the *decision weights* that map probabilities into the weights with which they appear in the evaluation of the prospect.

To allow addressing the broader question of what kind of incentive contract the principal might prefer—as opposed to optimizing within a specific class of contracts (as, for instance, in an optimal contest design problem)—we use a general model for contracts encompassing a broad variety of incentive schemes, including those observed on major crowdsourcing platforms such as fixed prices to a fixed number of workers (MTurk), base payments with bonuses (oDesk), and open-entry tournaments with prizes to a small number of contestants (Kaggle, Topcoder).

2.1 Optimal contracts: Expected utility maximizers

We first investigate the nature of the optimal contract that a principal should use when the population of workers behaves according to the standard model of expected utility theory—namely, agents with decision weights that equal probabilities, $\pi(p) = p$, and risk-neutral or risk-averse preferences corresponding to concave functions $u(x)$ (risk neutrality corresponds to the special case of linear u). We show that the optimal contract, when agents are expected utility maximizers, always takes the form of a fixed price mechanism where the principal pays an optimally chosen number of agents a fixed price[2] independent of output quality (Theorem 3.1 in Easley and Ghosh [2015]). The proof hinges on agents' not being risk-seeking, and proceeds by arguing that for *any* total expected payout W the principal might make, that payout W incentivizes the highest participation when it is disbursed as fixed payments rather than via any output-contingent scheme.

2.2 Contract design: Prospect theory preferences

We then ask whether a principal would ever choose a fundamentally different contract for agents with prospect theory preferences than that for expected utility maximizing agents. To address this question, we compare the incentives created by contests, which are widely used in a variety of crowdsourcing applications, against

[1]While the central issue in most of the principal-agent literature centers around effort elicitation—providing the right incentives for workers to elicit high effort (as a means to high output quality) in situations with imperfectly observable effort, incentives for *participation*, rather than quality, are the central issue in several online labor and crowdsourcing settings. This is due to differences in both the nature of available monitoring tools (various online labor platforms (including oDesk) offer virtual office applications for monitoring of workers) as well as intrinsic motivation for the task at hand: a pattern discovered by multiple experimental studies investigating worker behavior in online crowdsourcing platforms is that a change in the level of offered incentives affects participation rates, but not the degree of effort that workers exert conditional on participation (Mason and Watts [2009], Jeon et al. [2010], Kraut and Resnick [2012]).

[2]This price is c; if the value of c is private information to agents it can be elicited using an auction mechanism

the fixed payment contracts which are optimal for expected utility maximizing agents: for a given population of agents with preferences described by some value and decision weight functions, are there any values of the total amount the principal wants to spend on his task, W, for which he is better off using a winner-take-all contest with prize W instead of a fixed-payment contract with total payment W?

Contests are a very different kind of contract than fixed-payment schemes, in a qualitative sense: in contrast with fixed-payment schemes, an agent's reward in a contest is output-contingent, depending not only on her own output quality q_i but also on the draws of all other participants in the contest. So when might a contest, which is inherently a riskier prospect, provide stronger incentives for participation than the fixed-payment contract with its output-independent payoffs?

Contests and prospect theory preferences. If agents used actual outcome probabilities as their decision weights, Theorem 3.1 in Easley and Ghosh [2015] can be extended to show that a principal would never use output-contingent payments for any risk-neutral or risk-averse population of agents, no matter what the nature of his task. With prospect theory preferences, however, a small chance of winning a large prize might contribute more than its 'true' share of utility due to the overweighting of small probabilities, potentially creating a larger perceived payoff for the same expected payout to the principal.

To understand the intuition behind why contests might provide stronger incentives with overweighting of small probabilities, assume for a moment that an agent will 'win' the prize W if her quality draw beats an exogenous threshold q^* (imagine, for example, that the principal derives a value W if the output quality is greater equal the current state of the art q^*, and no value otherwise). This event $q_i \geq q^*$ has probability $\epsilon = 1 - F(q^*)$, and the perceived payoff to the agent from the contest is $u(W)\pi(\epsilon)$. A different payment scheme that has the same expected payout to the principal is the following: rather than wait to see if the output quality exceeds q^* before making the payment, the (risk-neutral) principal commits to paying his expected value, which is $W\epsilon$, yielding payoff $u(W\epsilon)$ to the worker. To develop intuition, imagine that the agent is risk-neutral so that $u(x) = \alpha x$; the agent then perceives respective payoffs $\alpha W \pi(\epsilon)$ and $\alpha W \epsilon$ from the output-contingent contest and fixed-payment schemes respectively. Now if the chance of winning the contest $\epsilon = 1 - F(q^*)$ is small enough for the overweighting of small probabilities to become relevant, the agent overweights the outcome of winning W ($\pi(\epsilon) > \epsilon$), leading to a larger perceived benefit ($\alpha W \pi(\epsilon)$) from the contest than from the fixed-payment scheme ($\alpha W \epsilon$), and correspondingly stronger incentives for participation. Of course, this reasoning is oversimplified—it ignores risk aversion, corresponding to non-linear u; more importantly, it ignores the fact that the threshold q^* that an agent's output q_i needs to beat in a contest is not exogenous but rather is endogenously determined in equilibrium by the choices of all other workers: $q^*(N^*) = \max_{j=2,\ldots,N^*}\{q_j\}$ where N^* is the equilibrium number of participants in the contest. However, it illustrates why overweighting of small probabilities can skew incentives in favor of 'riskier' gambles over 'safer' prospects, under the right circumstances.[3]

[3]This same reasoning can be used to understand why an output-dependent bonus payment, while

Despite the overweighting of small probabilities, however, a contest nonetheless does *not* always dominate a fixed-payment scheme: if the prize from winning the contest is not much larger than the cost of undertaking the associated task, few enough agents participate in equilibrium so that the probability of winning the prize becomes too large to be overweighted. This occurs even with risk-neutral agents; when agents display risk and loss aversion, the overweighting of probabilities—favoring the contest—needs to also compensate for this aversion to uncertainty if a contest is to dominate a fixed-payment contract.

Comparing contests against fixed-payment contracts. We first derive a necessary and sufficient condition under which a contest will dominate the fixed-payment contract for general preferences (u, π). Our first application of this result shows when this condition *cannot* be satisfied: we show that, for any prospect theory preferences, a principal who does not have a large enough budget to spend on the task should not conduct a contest, even for agents who overweight small probabilities—for such small prizes, risk and loss aversion beat out the benefits from overweighting of small probabilities.

More specific structure on the preference functions is required to say when the condition for a contest to dominate a fixed-payment contract *can* be satisfied. For this, we use the functional forms for u and π from the literature on econometric estimation of the cumulative prospect theory model [Tversky and Kahneman 1992; Bruhin et al. 2010]. Since these functional forms are derived from fitting the prospect theory model to extensive experimental data, they yield the best answer we can hope to have to our question—without conducting a measurement of the functions u and π in the specific marketplace of interest—for 'real' population preferences u, π.

Our main results, Theorems 4.4 and 4.5 in Easley and Ghosh [2015], together with the estimated parameter values from [Tversky and Kahneman 1992; Bruhin et al. 2010], provide an affirmative answer to our central question: we find that if the parameter describing the degree of risk aversion in u is larger than the parameter describing the degree of probability weighting in π (as is the case in the estimated model), a contest will eventually dominate a fixed payment scheme for large enough total payouts W. To the extent that these parameters indeed describe the decision-making behavior of agents in online crowdsourcing environments, therefore, a principal who values the output from crowdsourcing his task sufficiently highly (compared to the cost to a worker to produce that output) might indeed choose a different kind of mechanism, and do better as a result, if he accounts for deviations from the standard economic model of expected utilities that are typically used in theoretical design.

3. CONCLUSION AND FURTHER DIRECTIONS

Easley and Ghosh [2015] explores the idea that behavioral biases can indeed make a fundamental difference to a principal's choice of mechanism, via the problem of

strictly dominated for risk-averse agents with linear weighting of probabilities, may similarly improve incentives if agents overweight small probabilities for the same expected payout to the principal.

optimal contract structure in online crowdsourcing markets with participation-only strategic choices. While our model is a reasonable description of the specific crowdsourcing markets we use to motivate it, it is deliberately chosen to be the simplest, most parsimonious possible model that illustrates this idea—that deviations from expected utility theory, according to prospect theory preferences, are potentially quite significant to theoretical design—in a realistic setting. This means that there are a number of complexities of crowdsourcing markets that our model does not capture, presenting promising, and important, directions for further work, most interesting amongst which is when agents have heterogeneous preferences.

More generally, however, the question of how choice according to prospect theory versus expected utility affects equilibrium analysis and optimal design extends to domains beyond principal-agent problems in online labor: principal-agent problems are merely one amongst many applications of mechanism design where agents make decisions under uncertainty. A number of other domains to which mechanism design applies, especially those motivated by online systems where individuals rather than large firms are decision makers—such as online auctions with small bidders, or reputation systems in peer-to-peer economies—depend on the decisions of individual agents to whom these behavioral biases do apply. The issue of whether agents choose according to prospect theory versus expected utility preferences is a fundamental component of equilibrium analysis and the optimal design question in each of these environments, and as such, lays open a vast field of problems for further exploration.

REFERENCES

Bruhin, A., Fehr-Duda, H., and Epper, T. 2010. Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica 78,* 4.

Easley, D. and Ghosh, A. 2015. Behavioral mechanism design: Optimal crowdsourcing contracts and prospect theory. In *Proceedings of the Sixteenth ACM Conference on Electronic Commerce*. EC '15.

Ghosh, A. and Kleinberg, R. 2014. Behavioral mechanism design: Optimal contests for simple agents. In *Proc. 15th ACM Conference on Economics and Computation (EC)*. 913–930.

Jeon, G. Y., Kim, Y.-M., and Chen, Y. 2010. Re-examining price as a predictor of answer quality in an online Q&A site. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10.

Kahneman, D. and Tversky, A. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica 47,* 2.

Kraut, R. and Resnick, P. 2012. *Building successful online communities: Evidence-based social design*. Cambridge University Press.

Mason, W. and Watts, D. 2009. Financial incentives and the "performance of crowds". In *Proceedings of the ACM SIGKDD Workshop on Human Computation*. HCOMP '09.

Neumann, J. V. and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.

Savage, L. 1954. *The Foundations of Statistics*. Wiley.

Tversky, A. and Kahneman, D. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty 5*.