Eliciting Truthful Unverifiable Information

Extended Abstract

Shani Alkoby
The University of Texas at Austin
Austin, Texas
shani.alkoby@gmail.com

David Sarne
Bar-Ilan University
Ramat-Gan, Israel
david.sarne@gmail.com

Erel Segal-Halevi
Ariel university
Ariel, Israel
erelsgl@gmail.com

Tomer Sharbaf
Israel Ministry of Finance
Jerusalem, Israel
tomersrb@gmail.com

ABSTRACT

In many situations, an uninformed agent (UA) needs to elicit information from an informed agent (IA), typically when the latter has some unique expertise or knowledge related to some opportunity available to the UA. In many of those situations, the correctness of the information cannot be verified by the UA, and therefore it is important to guarantee that the information-elicitation mechanism incentivizes the IA to report truthfully. This paper presents and studies several information-elicitation mechanisms that guarantee truthful reporting, differing in the type of costs the IA incurs in producing and delivering the information. We show that with no such costs truthful information elicitation is possible with a positive but arbitrarily small expense for the UA. When information-delivery is costly, truthful information elicitation is possible where the extra expense for the UA (above the unavoidable cost of delivery) is arbitrarily small. Finally, when the information-production is costly, under some realistic condition related to the ratio between the expected gain of the IA from true reporting and the information-production cost, truthful information elicitation is possible where the extra expense for the UA (above the unavoidable cost of production) is arbitrarily small.

Full version is available at arXiv under the name "How to Make an Appraiser Work for You", http://arxiv.org/abs/1804.08314.

KEYWORDS

Information asymmetry; Information elicitation; Information disclosure; Unverifiable information

ACM Reference Format:
Shani Alkoby, David Sarne, Erel Segal-Halevi, and Tomer Sharbaf. 2018. Eliciting Truthful Unverifiable Information. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

In many real-life situations we need to know the value of some objects or opportunities available to us, but do not have the expertise to calculate it ourselves. Here are some examples. (i) You find a dazzling jewelery collection in a locked suitcase hidden in your attic and want to sell it. Knowing the exact value of each jewel can be very helpful for you when negotiating with potential buyers; (ii) You are a government owning some oil fields. You need to estimate the potential revenue from each field in order to know which of them is worth developing [6, 7, 9]; (iii) You sell a used car. You need to know its value in order to decide how much to ask for it.

A common solution in these situations is to buy the desired information from an expert [2–4, 8, 10], such as a jewel appraiser in 1, a geologist in 2 or a car mechanic in 3. Henceforth we call such an expert *the IA* (Informed Agent). The problem is that, in many cases, the information is not verifiable at the time it is delivered [5]. When you pay the IA, you have no way of knowing whether the provided information is correct [1]. Perhaps you will realize the true value later on (e.g, when you actually dig for oil), but at that later time, the IA might already be far away with your money in her pockets or the loss due to using the wrong value is already irreversible.

The problem of eliciting the true value becomes even more challenging when calculating and delivering the true value is costly for the IA, e.g., it requires the IA to invest time and effort. Here, the IA has a strong incentive to avoid calculating the true value and report an arbitrary value, since the Uninformed Agent (the UA) cannot tell the difference. The goal of the present research is thus to —

> *— develop mechanisms that incentivize an informed agent to calculate and deliver* true *information, even when calculation or delivery is* costly*, and even when the information is* unverifiable *by the uninformed agent.*

In the paper we first present the information-provision setting in a formal way. Then, we introduce mechanisms for truthful information elicitation for three types of settings, differing in the level of effort required for providing the information by the IA.

In the first setting, the IA incurs no cost for producing or delivering the information, hence it is used as a baseline. Here, our mechanism elicits the exact value of each opportunity. The expected expense of the UA is positive but arbitrarily small.

In the second setting, the IA incurs a cost for delivering the information to the UA. With some adaptations, our mechanism for the baseline case can still elicit the exact value of each opportunity. The expected expense of the UA in this case is quite favorable, as it only adds an arbitrarily small positive amount to the delivery cost. The latter is unavoidable in any mechanism, as the IA is self-interested and unless covers her information delivery cost will not

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, I. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

be willing to provide the information. For the second setting we provide two alternative mechanisms that elicit less accurate (yet still truthful) information for a potentially smaller expense.

In the third setting, the IA incurs a cost for producing the information. I.e, the IA initially does not have the information, but can obtain it by a costly effort. Here, the challenge is to convince the IA to both produce and provide the correct information, rather than just providing an arbitrary value. To handle this setting we make the realistic assumption that the cost of producing the information is sufficiently small relative to the expected value of the opportunity. Under this assumption, our mechanism is truthful and adds only an arbitrarily small amount to the unavoidable expense, which is the cost of producing the information.

2 THE MODEL

We consider a two-agents model in which one agent, the IA (Informed Agent), holds some information pertaining to *opportunities* or *objects* $O = \{o_1, ..., o_n\}$ that are available to other agent, the UA (Uninformed Agent), who does not have access to this information. By having the information related to an opportunity $o_i \in O$ that is available to an agent, the agent can make the most informed decisions related to that opportunity, e.g., whether it should be exploited or not, resulting in some profit (or loss). Therefore we can map each possible information to be received from the IA, assuming truthful, to the corresponding worth of the opportunity to the agent if both the opportunity and the information are available to the agent. We assume that the IA can transfer the information related to any opportunity $o_i \in O$ (and consequently its corresponding worth) to the UA. For exposition purposes and WLOG we assume that the information wanted by the UA regarding an opportunity o_i is fully encoded in the value of that opportunity, denoted by v_i. Each v_i is a random variable with probability distribution function $f_i(v)$ defined over the interval $[v^{\min}, v^{\max}]$, with $v^{\min} < v^{\max}$ such that $v^{\max} > 0$ and $v^{\min}$ is either positive or negative. These distributions are common knowledge, and there is no dependency between the values of two different opportunities. The corresponding cumulative distribution function of any probability distribution function $f_i(v)$ is denoted by $F_i(v)$.

The model assumes the UA cares about exact information only if the value of the opportunity is positive. This is because, if the value of an opportunity is negative, the UA is not going to exploit that opportunity anyway. On the other hand receiving the exact information in case the value is positive is crucial as the information prescribes how to generate such value. Formally, denote $v^+ := \max(v, 0)$. Then, for the UA, knowing the value v_i^+ is considered "exact", or "truthful". Hence, if the value of the opportunity is $v \leq 0$ then any report of a value $v' \leq 0$ is considered truthful.

The model assumes that any information transferred from the IA to the UA is practically unverifiable, i.e, either requires too much resources to verify, or can be verified only ex-post (after making an irreversible decision related to the opportunity). The strategy space available to the UA for eliciting information from the IA includes offering and/or requesting a payment to/from the IA as well as awarding the IA a subset of O such that the latter can benefit from them. The model assumes that the agents have quasi-linear preferences, so that the utility of the IA from getting a certain opportunity for a certain price is the opportunity's value (or zero if the value is negative) minus the price paid. The goal of the UA is to elicit truthful information from the IA while minimizing his own expected expenses, defined as the net payments made and the value of the opportunities transferred to the IA. The goal of the IA is to maximize her net expected profit, taking into account any costs incurred, payments received and values of awarded opportunities.

3 OVERVIEW

We present a short overview of the three settings mentioned above.

3.1 Case 0: No costs

We begin with a baseline case where the IA incurs no cost in producing and delivering the information. In this baseline case, we can extract the exact value of each opportunity by putting it to a variant of a Vickrey auction with a random reserve price. Moreover, the auction remains truthful even if it is done with an arbitrarily low probability. Therefore the UA's expense is arbitrarily small. Interestingly, and in contrast to the following cases, when information extraction and delivery is not associated with any cost, the mechanism for eliciting truthful information does not depend on the underlying distributions of values ($f_i(x)$ for $o_i \in O$).

3.2 Case A: Costly information delivery

In this section we assume that the IA already knows the true values of the opportunities, but disclosing these values to the UA incurs a publicly known cost. The mechanism of Section 3.1 can be augmented to elicit truthful reporting by simply adding the cost of delivering the set of opportunities in which the UA is interested. This puts the agents back in the situation of Case 0. Then, the mechanism of Section 3.1 can be used as-is to elicit the exact value of each opportunity in the wanted opportunities set.

However, in contrast to Case 0, now eliciting the exact values of all opportunities is not always the optimal action. We present two alternative mechanisms, the first allows us to elicit all opportunities with value above a threshold and the second allows us to elicit the k highest-valued opportunities. We numerically illustrate a situation in which the two less accurate mechanisms substantially improve the net gain of the UA relative to the exact-values mechanism.

3.3 Case B: Costly information production

In this section we assume that the IA does not know the true values, but can calculate them by incurring a publicly known cost. In contrast to Case A, here it will not help to just give the cost of production to the IA, since in this case, the IA will just take the money and provide some random value without actually calculating the true value. Our solution is to modify the mechanism of Case 0 by selecting the random reserve-price according to a carefully-designed distribution. We show that the mechanism is truthful as long as the following condition is satisfied:

$$c_i < \Pr\left[v_i > E[v_i]\right] \cdot E\left[v_i - E[v_i] \middle| v_i > E[v_i]\right] \quad (1)$$

where c_i is the cost of calculating the value of opportunity i, and Pr and E are the probability and expectation operators determined by the prior distribution of the value v_i. In case this condition holds, the UA's expense (above the production cost c_i) is arbitrarily small.

REFERENCES

[1] George A Akerlof. 1970. The market for "lemons": Quality uncertainty and the market mechanism. *The quarterly journal of economics* (1970), 488–500.
[2] Shani Alkoby and David Sarne. 2017. The Benefit in Free Information Disclosure When Selling Information to People.. In *AAAI*. 985–992.
[3] Shani Alkoby, David Sarne, and Sanmay Das. 2015. Strategic free information disclosure for search-based information platforms. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 635–643.
[4] Shani Alkoby, David Sarne, and Igal Milchtaich. 2017. Strategic Signaling and Free Information Disclosure in Auctions.. In *AAAI*. 319–327.
[5] Moshe Babaioff, Robert Kleinberg, and Renato Paes Leme. 2012. Optimal mechanisms for selling information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 92–109.
[6] Kenneth Hendricks, Robert H Porter, and Guofu Tan. 1993. Optimal selling strategies for oil and gas leases with an informed buyer. *The American Economic Review* 83, 2 (1993), 234–239.
[7] Kenneth Hendricks, Robert H Porter, and Charles A Wilson. 1994. Auctions for oil and gas leases with an informed bidder and a random reservation price. *Econometrica: Journal of the Econometric Society* (1994), 1415–1444.
[8] Benny Moldovanu and Manfred Tietzel. 1998. Goethe's SecondâĂŘPrice Auction. *Journal of Political Economy* 106, 4 (1 Aug. 1998), 854–859. https://doi.org/10.1086/250032
[9] Robert H Porter. 1995. The Role of Information in US Offshore Oil and Gas Lease Auction. *Econometrica: Journal of the Econometric Society* (1995), 1–27.
[10] David Sarne, Shani Alkoby, and Esther David. 2014. On the choice of obtaining and disclosing the common value in auctions. *Artificial Intelligence* 215 (2014), 24–54.