

The Benefit in Free Information Disclosure When Selling Information to People

Shani Alkoby

Bar-Ilan University, Israel
shani.alkoby@gmail.com

David Sarne

Bar-Ilan University, Israel
sarned@cs.biu.ac.il

Abstract

This paper studies the benefit for information providers in free public information disclosure in settings where the prospective information buyers are people. The underlying model, which applies to numerous real-life situations, considers a standard decision making setting where the decision maker is uncertain about the outcomes of her decision. The information provider can fully disambiguate this uncertainty and wish to maximize her profit from selling such information. We use a series of AMT-based experiments with people to test the benefit for the information provider from reducing some of the uncertainty associated with the decision maker's problem, for free. Free information disclosure of this kind can be proved to be ineffective when the buyer is a fully-rational agent. Yet, when it comes to people we manage to demonstrate that a substantial improvement in the information provider's profit can be achieved with such an approach. The analysis of the results reveals that the primary reason for this phenomena is people's failure to consider the strategic nature of the interaction with the information provider. People's inability to properly calculate the value of information is found to be secondary in its influence.

Introduction

Information providers have become an integral part of almost every aspect of modern life. Information providing can have many forms. For example, an information provider can be an expert: a weatherman offering a forecast for the weather on our upcoming wedding day, an accountant offering information about the worth of a company we plan to take over on or a mechanic offering a reliable estimate for the value of a used car we intend to buy. Another example is information platforms or services such as TripAdvisor that provides traveling-related information, CarFax that provides detailed reports on used cars and online credit-report services. Common to all the above, that they aim to provide the user with information about the identity, nature or value of the different opportunities that can potentially become available to her in a given environment. The information provided eliminates much of the uncertainty associated with the outcomes of different decision alternatives, hence its value.

While the study of information providers in multi-agent settings is extensive, the focus of prior work aiming to

study strategic behavior of such entities is mostly limited to how they should price their services (Alkoby, Sarne, and David 2014; Lai et al. 2014; Cheng and Koehler 2003; Hajaj and Sarne 2014; Sarne, Alkoby, and David 2014b). In this paper we investigate a complimentary means aiming to increase the likelihood of information purchase—an a priori free information disclosure.

The main idea of free information disclosure is that through the elimination of some of the possible outcomes, knowing the true outcome becomes highly valuable. For example, consider a passenger that is about to go on a flight from NY to Paris in order to attend an important business meeting. Now suppose the possible outcomes of the flight are: (i) arriving on time, with an a priori probability of 94.4%; (ii) arriving an hour late, with an a priori probability of 4.1%; or (iii) missing the meeting because the flight gets canceled due to a union strike, with an a priori probability of 1.5%. Knowing the true outcome (e.g., by purchasing it from an oracle or a corrupted union member) has very little value, as the chance of not arriving to the meeting on time is very small (1.5%). However assume the oracle publicly announces that the flight is not going to arrive on time (i.e., eliminating the first outcome, hence reducing the set of possible outcomes to the latter two). Now, there is much value in being able to distinguish between the two remaining outcomes—the naive posterior probability of ending up with a canceled flight due to a strike is 27%. Therefore the passenger will be willing to pay a substantial amount in order to obtain this information. Still, the above naive probability update process does not take into account the strategic considerations that lead the information provider to publicly disclose some of the information she holds. The incorporation of the strategic aspect of the interaction results in a somehow different probabilistic update and in fact we can prove that free information disclosure is necessarily detrimental in this case. However, when dealing with people, the above does not necessarily hold. It is well known that people are often irrational (Rabin 1998; Kahneman 2000; Azaria et al. 2015; Buntain, Azaria, and Kraus 2014). Therefore, it is possible that they will not take into consideration the strategic nature of the interaction or even fail to properly reason about the value of information to some extent, making free information disclosure beneficial for the information provider.

This paper provides a comprehensive experimental evalu-

ation of the above approach whenever interacting with people, attempting to identify the main sources of people’s inability to make the right decision when it comes to information purchasing. It uses a testbed that captures a core “value of information” problem setting of the kind described above, with an information provider that can fully disambiguate the uncertainty associated with outcomes. The experiments involve 300 subjects experiencing a total of 6000 information purchasing decisions, interacted through Amazon Mechanical Turk using three different treatments.

Contributions. The paper makes two main contributions. The first is showing that, unlike with fully rational buyers, when it comes to people free information disclosure can substantially improve the information provider’s profit from selling information. The second is showing that the improvement achieved is mostly because of people’s inability to take into consideration the strategic nature of the interaction rather than their somehow limited ability to properly calculate the value of information.

The Model

We consider the basic standard model of a self-interested information provider and a prospective information buyer (denoted “buyer” onwards). The buyer is facing a simple decision problem involving an opportunity O available to her, where the possible available alternatives are to exploit opportunity O or opt-out not to exploit it. The set of possible exploitation outcomes (corresponding to different possible nature states) is denoted $V = \{v_1, v_2, \dots, v_n\}$, where the corresponding a priori probability of each value $v \in V$ is captured by the function $p(v)$ ($\sum p(v_i) = 1$). If choosing to opt-out the buyer gains some fallback profit v_0 . The buyer and the information provider are symmetric in the sense that they are both familiar with V and the function $p(v)$. The information provider is also acquainted with the true state of the world, i.e., knows the true exploitation value of O and can sell this information to the buyer for a fee. In an effort to increase her profit, the information provider can use a strategic behavior and publicly eliminate some of the possible outcomes of O such that this information becomes available to the buyer before she makes her decision of whether to purchase the identity of the true outcome or not. We denote this latter strategy PFID (Preliminary partial Free Information Disclosure) for short.

The course of the game is therefore as follows: nature first sets the true exploitation outcome v of the opportunity O ; the value v becomes available to the information provider who sets the requested fee c for revealing v along with eliminating (publicly) some of the values in V , such that the remaining possible values are those in the subset $D \subseteq V$; based on c and D , the buyer can decide either to purchase the true exploitation value of O , in which case the value v is disclosed to her, or not; finally, the buyer decides whether to exploit opportunity O .

The model assumes that both the exploitation values and the cost of purchasing the information from the information provider are additive. The goal of the buyer is therefore

to maximize her expected profit, defined as the exploitation value of O (if exploiting the opportunity) or the fallback v_0 (otherwise) minus the payment to the information provider (if purchasing the information).

The above model can be mapped to various real-life problems. For example, the buyer can represent a company that considers taking over its competitor. The true value of the other company is uncertain however can be purchased from an internal source that may increase the value of the information it holds through PFID. The information provider’s problem is thus, given an opportunity O and the true exploitation value v , which exploitation values to eliminate for free and what price to set as the fee for revealing v in order to maximize her expected profit.

Rational Buyers

We first analyze the best response strategies and the resulting equilibrium in case the buyer is fully rational and risk neutral.

Buyer

In the absence of any preliminary information from the information provider, the buyer will choose to exploit opportunity O only if the expected exploitation value is greater than the fallback value v_0 . The buyer’s **Expected Monetary Value (EMV)** is thus given by:

$$EMV(O) = \max\left(\sum_{v \in V} v \cdot p(v), v_0\right) \quad (1)$$

If purchasing the information from the information provider, the buyer’s decision is made under certainty. Here, the buyer exploits O only in cases where the exploitation value is greater than the fallback utility. The **Expected Value Under Certainty (EVUC)** is thus given by:

$$EVUC(O) = \sum_{v \in V} \max(v, v_0) \cdot p(v) \quad (2)$$

The **Value of the Information** held by the information provider to the buyer (denoted VoI onwards) is thus $VoI(O) = EVUC(O) - EMV(O)$ and this is the maximum amount the buyer will be willing to pay for receiving the true outcome.

Finally, when the information provider uses PFID, leaving only a subset D of remaining applicable outcomes, the above calculations still hold with some minor modifications:

$$VoI(D) = EVUC(D) - EMV(D) = \sum_{v \in D} \max(v, v_0) \cdot Pr(v|D) - \max\left(\sum_{v \in D} v \cdot Pr(v|D), v_0\right) \quad (3)$$

where $Pr(v|D)$ is the posterior probability of the exploitation value being v given the evidence D . Naively, the value of $Pr(v|D)$ should be calculated through a simple update of the a priori probability $p(v)$ as follows:

$$Pr(v|D) = \begin{cases} \frac{p(v)}{\sum_{y \in D} p(y)} & \text{if } v \in D \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

The above calculation is considered naive as it does not take into consideration the strategic behavior of the information

provider. Recall that the information provider's strategy is a function $S : V \rightarrow D$ (and the corresponded prices to be charged for revealing the true value, calculated as the VoI), specifying for each outcome $v \in V$ the subset $D \subseteq V$ of remaining possible exploitation values. The strategy thus induces a partition of the set V such that any two outcomes v_i and v_j are in the same partition element if and only if $S(v_i) = S(v_j)$. Since in equilibrium the buyer is using her best response strategy to the information provider's strategy, the posterior probabilities calculation taking place by the buyer should be based on S and is given by:

$$Pr(v|D) = \begin{cases} \frac{p(v)}{\sum_{y|S(y)=D} p(y)} & \text{if } S(v) = D \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

Information Provider

As explained above, the information provider sets the cost of her information providing service to VoI. The information provider may attempt to maximize the VoI through PFID. In this case we distinguish between having a naive buyer and a strategic one.

Naive Buyer When the buyer does not take into consideration the fact that the information provider is acting strategically she uses the naive probability update according to (4). For example, assume that $V = \{-100, 0, 100\}$ where all values are possible with equal probability, and assume the fallback if not exploiting the opportunity is $v_0 = 0$. Here $\text{VoI} = 33.3$. However, if the information provider uses $S(-100) = S(100) = \{-100, 100\}$ and $S(0) = \{-100, 0, 100\}$, she can still charge 33.3 whenever $v = 0$ however charge 50 in case $v \in \{-100, 100\}$. Generally, when facing a naive buyer, the information provider should choose for every value $v \in V$ to eliminate all exploitation values except those in $D \subseteq V$ such that the difference $EVUC(D) - EMV(D)$ is maximized and charge exactly the difference.

Strategic Buyer When the buyer is acting strategically, however, the information provider cannot benefit from free information disclosure, as stated in Proposition 1.

Proposition 1. *The information provider's expected profit when using PFID is bounded (from above) by the expected profit when not using it.*

Proof. Since the information provider sets the price of her service according the worth of the information, we need to show that the expected VoI under certainty without PFID is greater than with PFID. Meaning that the following holds:

$$EVUC(O) - EMV(O) \geq \sum_D (EVUC(D) - EMV(D)) \cdot Pr(D) \quad (6)$$

where $Pr(D)$ is the probability the buyer will receive the information D , calculated according to: $Pr(D) = \sum_{v \in D} p(v)$. Notice that $EVUC(O) = \sum_{v \in V} \max(v, v_0) \cdot p(v) = \sum_D Pr(D) \sum_{v \in D} \max(v, v_0) \cdot Pr(v|D) = \sum_D EVUC(D) \cdot$

$Pr(D)$. Therefore, in order for (6) to hold we only need to prove that $\sum_D EMV(D) \cdot Pr(D) \geq EMV(O)$:

$$\begin{aligned} EMV(O) &= \max\left(\sum_{v \in V} v \cdot p(v), v_0\right) = & (7) \\ \max\left(\sum_D Pr(D) \cdot \sum_{v \in D} v \cdot Pr(v|D), v_0\right) &\leq \\ \sum_D Pr(D) \cdot \max\left(\sum_{v \in D} v \cdot Pr(v|D), v_0\right) &= \\ \sum_D EMV(D) \cdot Pr(D) &\quad \square \end{aligned}$$

Therefore, if both the information provider and the buyer are fully rational and strategic, there is no point for the information provider to use PFID. In the following section, however, we show experimentally that there is much value in such strategy when the buyer is a person.

Irrational Buyers

In most real-world settings we expect to find people in the role of the buyer. This section describes an experiment carried out for testing the effectiveness of PFID in such a case.

Possible Failures in Decision Making

Prior work provides much evidence for people's bounded rationality in decision making situations in the sense that they do not adhere to rigid models of rationality and are easily influenced by various external factors and biased towards certain conclusions (Simon 1972; Levy and Sarne 2016; Kahneman 2000; Hajaj, Hazon, and Sarne 2015; Kraus 2015). Specifically, for the strategic interaction settings considered in this paper we identify two possible causes for irrational behavior that may affect the decision whether to purchase the information offered. The first is people's somewhat limited reasoning and computational capabilities that may prevent the proper calculation of the value encapsulated in the information according to the guidelines given in the former section. The second is people's failure to take into consideration the strategic nature of the interaction with the information provider. The implication of the latter is failure to update the probabilities assigned to the different exploitation values according to (5) and using the following naive calculation instead:

$$Pr(v|D) = \begin{cases} \frac{p(v)}{\sum_{y \in D} p(y)} & \text{if } S(v) = D \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

Meaning that the buyer does not take into consideration the reason the information provider decided to disclose D rather than any other subset.

We note that prior literature contains evidence for both above phenomena, i.e., people's failure to take into consideration the strategic aspect of an interaction (Eisenhardt and Zbaracki 1992) and failure to accurately calculate the value of information (Kamar, Gal, and Grosz 2013; Bazerman and Moore 2008). The extent of the effect, if any, depends on the domain, the nature of the interaction and the complexity of the underlying problem. Still, none of these works consider a model similar to ours and the results reported there cannot be trivially carried over to our case.

Among the two effects, the second clearly favors the use of PFID in a way that increases the value of the information held by the information provider whenever the buyer follows (4). The effect of the inability to properly calculate the value of information (i.e., even if taking into consideration the strategic aspect of the interaction) when using PFID is somehow vague, as it is not clear whether it will actually result in an increase or a decrease in the value buyers see based on the information provided, even in cases where $Vol(D)$ increases. Our experiments were designed such that both effects can be isolated to a great extent.

Experimental Framework

For our experiments we used a multi-round game called “What’s In The Box?”, which captures the essence of the basic underlying decision making problem in our model without adding any externalities that may confuse participants. On each round in the game the player is introduced with a box which contains a prize expressed in game points (corresponding to an opportunity in our model). The available alternatives are to open the box (corresponding to exploiting it) or leave it unopened (opt-out). Along with the box the player is also introduced with the possible values of the prize in it (corresponding to the possible exploitation values). Prize values can be either positive or negative, each having an a priori equal chance. Prior to her decision whether to open the box, the player can request to obtain the identity of the prize in the box, i.e., completely disambiguate the uncertainty associated with the value. This latter information is, however, costly, and the cost of obtaining it (expressed in terms of game points) is provided to the player prior to making her decision to request it. The player thus needs to decide whether to purchase the information about the true value of the prize in the box and then whether to open the box. If choosing not to open the box the player obtains zero game points (the fallback value). Finally, the player moves on to the next game round, and the appropriate adjustments to her total accumulated game points are made (adding the prize (or actually reducing it in case its value is negative) in case the box was opened and reducing the cost of information if purchased). The goal of the player is to accumulate as many game points as possible throughout the game.

We note that the primary reason for choosing a repeated game where on each round the player is facing a different decision problem instance (though of similar nature) was to have people follow an EMV-based decision rule. Prior work provides much evidence for the fact that in repeated-play settings people’s strategies asymptotically approach the EMV strategy as the number of repeated plays increases (Klos, Weber, and Weber 2005; Keren 1987; Barron 2003). The proper solution to the game, when taking an EMV-maximizing approach is quite straightforward and follows exactly the calculation given in the section dealing with rational buyers: the player should purchase the information if $EVUC(O) - EMV(O)$ (or $EVUC(D) - EMV(D)$ when using PFID) is greater than its cost, and open the box only if the value of the prize (or the expected value of the prize in case the information is not purchased) is greater than zero.

Experimental Design

We implemented the “What’s In The Box?” game using *C#.net* for the server side and *Html5, css* and *Javascript* for the client side such that participants could interact with the system using a relatively simple graphic interface. Figure 1 present a screen shot of the game where the player is introduced to the possible values of the prize in the box (all with the same probability) and the cost of purchasing the true value. In this stage, the buttons provided to the player to make her decision are disabled for the first ten seconds so she is forced to spend some time thinking before making her decision.

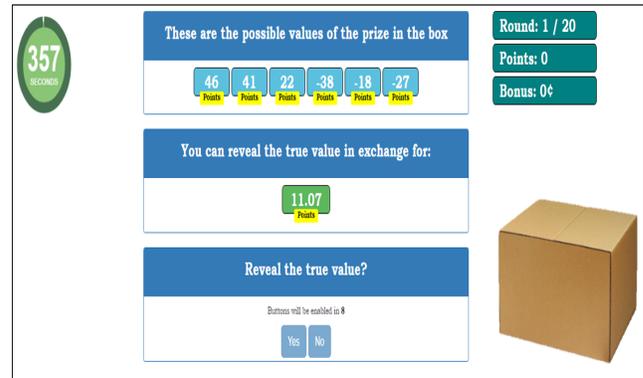


Figure 1: Screen shot of the game. See text for details.

In order to support free information disclosure, we enabled crossing out some of the possible values of the prize in the box few seconds after they appear, so that the player could still see the set of original values and those that have been removed. At the end of each round the player received a short summary detailing the change in her accumulated game points, listing the prize obtained (if opening the box) and the payment for the information (if purchased). We used three different experimental treatments:

No Free Information Disclosure - where no free information disclosure takes place, i.e., none of the values is crossed out prior to the information purchase decision.

Free Information Disclosure by an Explicitly Strategic Information Provider - where information is sold by a strategic information provider that uses PFID. With this treatment we did everything possible, from the UI point of view, to make sure the player understands that values are being eliminated by a self-interested agent that aims to maximize its own gain. Therefore the player was told that there is additional player in the game, who gains from selling the information to her. In each round, in addition to presenting the player’s own accumulated score on the screen, we also presented the information provider’s accumulated profit.

Free Information Disclosure by a Non Explicitly Strategic Information Provider - where information is sold by

a strategic information provider, except that with no mentioning of the strategic considerations accounting for the disclosure of information. Participants were told that values are removed by the “system” as a way of helping the player and obviously there was no mentioning or reflection of the information provider (or its score) in the GUI. The idea in including this treatment in our experiment was to see how close will be the decisions of players under this treatment to those exhibited in the second treatment. A great similarity would indicate that people tend to ignore the strategic nature of the information provider in the second treatment. The use of PFID in the last two treatments followed the guidelines provided at the end of the Naive buyer part in the Rational-buyers section above (i.e., maximizing the expected profit assuming facing a naive buyer).

In order to have better control over the experiment we pre-generated a set of core problem settings. The values for the different outcomes in each problem were integers randomly picked within the range $[-50, 50]$. In order to reason about the effect of the number of values on the results obtained we generated a total of 250 such problems, differing in their number of outcomes n , in a way that we had 50 problems for each number of outcomes $n \in [3, 7]$. In order to reason about the effect of the magnitude of the difference between the value of information and its cost on people’s ability to make the right decision, we took the cost to be exogenously set (rather than setting it as $\text{VoI}(O)$).¹ For this purpose we created four problem instances based on each core setting O (of the 250 mentioned) differing in the cost of purchasing the information, setting the price of information to: (1) $0.8 \cdot \text{VoI}(O)$; (2) $1.2 \cdot \text{VoI}(O)$; (3) $0.2 \cdot \text{VoI}(O)$; and (4) $1.8 \cdot \text{VoI}(O)$. In those few cases where $\text{VoI}(O)=0$ (e.g., when all outcomes are positive) we randomly picked the cost of information for each of the four resulting problem instances (within the range $[0, 50]$). The full set of problem instances is available upon request from the corresponding author. Overall, in 48% of the problems a rational buyer should purchase the information and in the remaining 52% she should not, where the difference corresponds to those cases where $\text{VoI}(O)=0$ (hence information should not be purchased regardless of its price).

Participants were recruited and interacted through Amazon Mechanical Turk (AMT) which has proven to be a well established method for data collection in tasks which require human intelligence (Paolacci, Chandler, and Ipeirotis 2010). To prevent any carryover effect a “between subjects” design was used, assigning each participant to one treatment only. The compensation for taking part in the experiment was composed of a show-up fee (the basic “HIT”) and also included a bonus, which was a direct outcome of the participant’s performances in the experiment (measured as the amount of accumulated game points), in order to encourage thoughtful participation—one cent bonus for each 10 game

points accumulated. Each participant received thorough instructions of the game rules, the compensation terms and her goal in the game. Then, participants were asked to engage in practice games until stating that they understood the game rules (with a strict requirement for playing at least two practice games). Prior to moving on to the actual games, participants had to correctly answer a short quiz, making sure they fully understand the game and the compensation method. Finally, participants were requested to play a sequence of 20 rounds, where the problem instance used for each round was randomly picked from the pool of 1000 problem instances described above (with no repetition).

During the game, we logged all player actions along the different phases (instructions, training, quiz and actual game). We had four classifications for each player’s information purchasing decision: whenever purchasing the information, the decision was classified as “good” if the VoI is greater than or equal to its cost (and “bad” otherwise). Similarly, whenever not purchasing, the decision was classified as “good” if the VoI is lower than or equal to its cost (and “bad” otherwise). The above was calculated in all three treatments according to the naive VoI calculation as described in (8). For the two treatments that use PFID, we repeated the calculation by taking the VoI to be calculated according to (5) assuming the information provider applies the PFID as described above.

Results

Overall, we had 300 participants taking part in our experiments, 100 for each experiment, each playing 20 rounds according to the above design. Participants ranged in age (18-81, average 34.5) and gender (57% men and 43% women), with a fairly balanced division between treatments.

The analysis of the results shows that our virtual information provider managed to substantially improve the overall profit from selling information when using PFID, compared to when not using it. The following table details the average per-game (20 rounds) profit obtained with each of the three treatments and the number of instances in which the player purchased the information.

	Treatment 1	Treatment 2	Treatment 3
Avg. Total Profit	57.2	73.6	77.4
# of sales	1030	1206	1189

The above table reflects an increase of 29% and 35% (both statistically significant using $t - test$, $p < 0.005$) in the information provider’s profit through PFID (compared to when not using it), when presenting the information provider as a fully strategic player (treatment 2) and when avoiding any mentioning of its strategic nature (treatment 3), respectively. The improvement in expected profit due to not presenting the information provider as a strategic player (i.e., in the transition from treatments 2 to 3) is 5% (non statistically significant using $t - test$, $p > 0.5$). Similarly, the number of instances in which the information provider managed to sell the information she was holding when using PFID (i.e., (treatments 2 and 3), increased by 17% and 15% (both sta-

¹As otherwise, if the information provider sets the price to be exactly $\text{VoI}(O)$ even the slightest deviation in the calculation of the value of information may lead to wrong results, precluding a genuine analysis of the extent to which people are affected by their failure to take the information provider as a strategic agent.

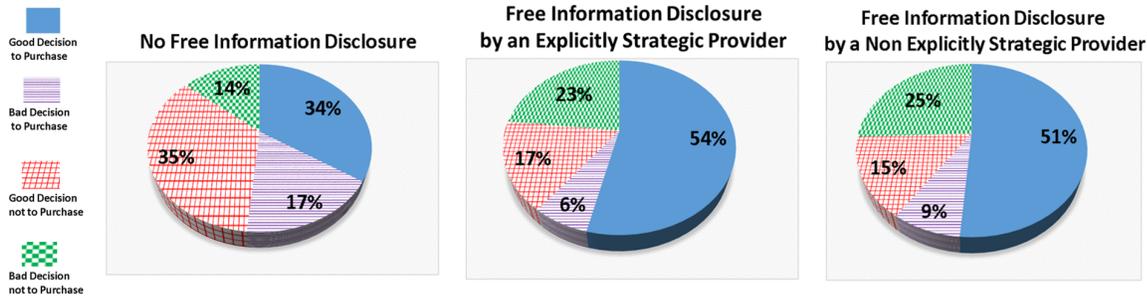


Figure 2: Classification (using naive VoI calculation) of decisions made in all treatments.

tistically significant using $t - test$, $p < 0.005$) with the two information-disclosure treatments (and a minor reduction of 1% in the transition in-between the last two (non statistically significant using $t - test$, $p > 0.5$)). The insignificant differences between the profits obtained with treatments 2 and 3, as well as further similarities observed in the in-depth analysis of the results, as reported in the following paragraphs, suggest that people do not take into consideration the strategic aspect of the problem they are facing in this domain. One additional evidence that strengthens this latter hypothesis can be found in the performance achieved in the third treatment. If the players were fully rational as far as the computation of the VoI is concerned, yet still naive in the sense of not taking the information provider to be strategic, then the theoretical expected profit of the information provider based on the 1000 problem instances is 85.76. The information provider in the third treatment, the one that emulates this exact scenario, managed to reach a very close profit (77.4).

Figure 2 provides a more detailed investigation concerning the sources of the improvement achieved with PFID. It depicts the break-down of the total 2000 information purchase decisions made in each treatment into the four different classifications described in the experimental design section (based on naive VoI calculation). Considering the chart that summarizes the results obtained when not using PFID (most left), we observe that the general success of people with the tested settings is 69%, with relatively similar chance of choosing the wrong action according to the two classifications (either purchase when better not to purchase and vice versa). These latter findings suggest that people are indeed unable to properly calculate the value of information to some extent. With the information providers using PFID we observe that the percentage of instances in which information is purchased increases from 51% to 60% (regardless of how the information provider was presented to the players). Interestingly, in the second and third treatments the percentage of cases where information was purchased out of those where it should not had been ($23/54 = 30\%$ and $25/76 = 33\%$, respectively) or when not purchased out of those where it should had been ($26\%, 37\%$) did not change much between treatments. This, as well as the relatively similar division into the four classifications observed within the charts corresponding to the second and third treatments, indicate, once again, that it is people's failure to consider the information

provider to be strategic that accounts for most of the improvement achieved in the information provider's profit. The inability to accurately calculate the value of information is definitely reflected in the results, as explained above, yet its impact is only secondary to the primary effect.

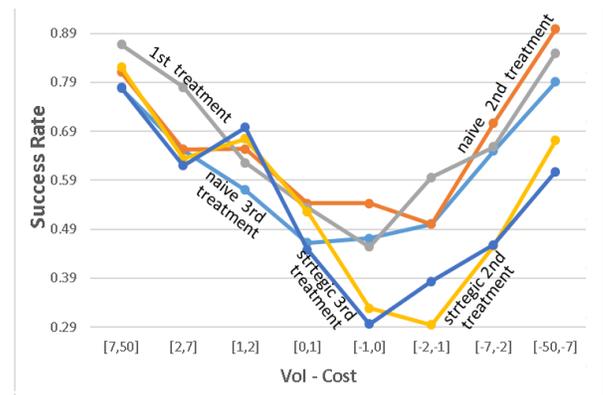


Figure 3: Players' success rate in the different treatments.

Figure 3 depicts the average success rate of players in their decision of whether to purchase the information from the information provider, as a function of the benefit in purchasing it (VoI-cost) in the different treatments. The success rate is measured as the percentage of decisions classified as "good" out of all those made. For each of the two treatments using PFID we included two curves. The first refers to classifications according to the naive calculation of the VoI and the second according to the calculation that takes all strategic considerations into account. From the graph we observe that people are quite good in realizing that purchasing the information is beneficial (or not beneficial) whenever the difference between the true value of the information and its cost is substantial. The greater the difference, the better the quality of the decision people make.

The fact that the two curves corresponding to the naive information value calculation under PFID almost entirely coincide with the curve corresponding to not using PFID suggests that people completely fail to consider the strategic behavior of the information provider. The improvement

achieved in the information provider's profit is thus primarily through the increase in the number of instances where the value of information becomes greater than its cost. Indeed, even with PFID people still reflect the same computational difficulties in reasoning about the benefit in purchasing the information, however since the overall number of "beneficial" instances increases so does the number of times information is purchased. The two curves representing the quality of people's decisions when the VoI calculation takes strategic considerations into account are very close. Their general behavior also reflects better success whenever the benefit in purchasing the information is relatively high or low, though their center point is shifted compared to the others. These two insights complement all the findings reported so far related to the role of the two hypothesized reasons in generating the benefit PFID achieves.

Related work

Decision making under uncertainty and reasoning about the value of information are prevalent themes in AI, commonly used in areas such as active learning, observation selection (also referred to as "set selection" and "prediction models for active learning" (Zhang 2010; Bilgic and Getoor 2011)) and user interaction (Kapoor, Horvitz, and Basu 2007; Krause and Guestrin 2009; Tolpin and Shimony 2012). Its main usage is as a sensitivity analysis technique to rate the usefulness of various information sources and to decide whether pieces of evidence are worth acquisition before actually using them (Liao and Ji 2008).

In human-computer interaction, much effort has been placed on modeling the user's attentional state in order to reason about the cost of (and consequently the benefit in) requesting information from the user or providing her with some information held by the system (Yakout et al. 2011; Horvitz et al. 2003; Hui and Boutilier 2006). While the underlying value of information calculation in these works is similar to ours, the information provider/requester they consider is fully cooperative in the sense that it attempts to maximize the user's expected benefit instead of its benefit from selling the information to the user as in our case.

Much work can be found in the multi-agent literature studying strategic information providers that can disambiguate the uncertainty associated with the opportunities available to agents (Sarne, Alkoby, and David 2014a; Moscarini and Smith 2003; Azoulay-Schwartz and Kraus 2004). These, however, primarily deal with the question of information pricing and do not incorporate the option for selectively disclosing some of the information in order to increase the chance for a purchase. Those that do consider the option to use selective information disclosure (in more complex decision settings, where the method can theoretically matter even when taking the strategic aspect of the interaction) (Alkoby, Sarne, and Das 2015; Hajaj and Sarne 2014), or even those that studied the role of information revelation (Dufwenberg and Gneezy 2002; Eső and Szentes 2007; Ganuza and Penalva 2010), typically assume the information consumers are fully rational agents.

The idea itself of selective information disclosure for affecting people's behavior is not new in general and can be

found in various other literature (Thaler and Sunstein 2008; Sheena 2010; Azaria et al. 2012; Hajaj, Hazon, and Sarne 2016; Peled, Kraus, and others 2015; Azaria et al. 2016). It has been justified in prior literature mainly as means for increasing user loyalty, attracting potential users, inducing repeated service requests or influencing the user's behavior (Rysman 2009; Elmalech, Sarne, and Grosz 2015). Nevertheless, to the best of our knowledge, an empirical investigation of the benefit in free information disclosure in order to promote information purchase by people has not been carried out to date.

Conclusions

The encouraging results reported in the results section suggest that information providers can benefit much from free information disclosure when facing human buyers. The importance of this finding is primarily due to the fact that real-world information buyers are human (as opposed to fully rational agents), the extensive penetration of strategic information providers to almost any field in our lives and the wide applicability of the underlying decision making model used. These results will be valuable both for practitioners developing information providing platforms and applications and for researchers who hopefully will see the potential in continuing this line of work and design and test more advanced methods for improving information providers' revenues when interacting with people, based on the insights provided in this paper.

The results' analysis unfolds the main reason for the success of the proposed approach: it is primarily people's failure to consider the strategic nature of the interaction that precludes a proper judgment. Therefore, an increase in the naive value of certainty translates to an almost identical increase in its value in the eyes of the buyer.

In future work we plan to add the price-setting capability into the information provider's strategy. In our experiments, the price was exogenously set for the reasons mentioned in the experimental design section. Having control both over the price asked and the information disclosed prior to the buyer's decision whether to purchase the information can certainly increase the information provider's profit. Still, this requires learning the mutual effects of these two parameters, as irrational buyers may be affected by different combinations of price and disclosed information in various ways.

Acknowledgments.

This research was partially supported by the ISRAEL SCIENCE FOUNDATION (grants No. 1083/13) and the ISF-NSFC joint research program (grant No. 2240/15). We are grateful to Anatoliy Tsvetov for his help in developing the experimental infrastructure.

References

Alkoby, S.; Sarne, D.; and Das, S. 2015. Strategic free information disclosure for search-based information platforms. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 635–643. International Foundation for Autonomous Agents and Multiagent Systems.

- Alkoby, S.; Sarne, D.; and David, E. 2014. Manipulating information providers access to information in auctions. In *Technologies and Applications of Artificial Intelligence*. Springer. 14–25.
- Azaria, A.; Rabinovich, Z.; Kraus, S.; Goldman, C. V.; and Tsimhoni, O. 2012. Giving advice to people in path selection problems. In *Proceedings of the AAMAS*, 459–466.
- Azaria, A.; Rabinovich, Z.; Goldman, C. V.; and Kraus, S. 2015. Strategic information disclosure to people with multiple alternatives. *ACM Transactions on Intelligent Systems and Technology (TIST)* 5(4):64.
- Azaria, A.; Gal, Y.; Kraus, S.; and Goldman, C. V. 2016. Strategic advice provision in repeated human-agent interactions. *Autonomous Agents and Multi-Agent Systems* 30(1):4–29.
- Azoulay-Schwartz, R., and Kraus, S. 2004. Stable repeated strategies for information exchange between two autonomous agents. *Artificial Intelligence* 154(1):43–93.
- Barron, B. 2003. When smart groups fail. *The journal of the learning sciences* 12(3):307–359.
- Bazerman, M. H., and Moore, D. A. 2008. Judgment in managerial decision making.
- Bilgic, M., and Getoor, L. 2011. Value of information lattice: Exploiting probabilistic independence for effective feature subset acquisition. *J. Artif. Intell. Res. (JAIR)* 41:69–95.
- Buntain, C.; Azaria, A.; and Kraus, S. 2014. Leveraging fee-based, imperfect advisors in human-agent games of trust. In *AAAI*, 916–922.
- Cheng, H. K., and Koehler, G. J. 2003. Optimal pricing policies of web-enabled application services. *Decision Support Systems* 35(3):259–272.
- Dufwenberg, M., and Gneezy, U. 2002. Information disclosure in auctions: an experiment. *Journal of Economic Behavior & Organization* 48(4):431–444.
- Eisenhardt, K. M., and Zbaracki, M. J. 1992. Strategic decision making. *Strategic management journal* 13(S2):17–37.
- Elmalech, A.; Sarne, D.; and Grosz, B. J. 2015. Problem restructuring for better decision making in recurring decision situations. *Autonomous Agents and Multi-Agent Systems* 29(1):1–39.
- Eső, P., and Szentes, B. 2007. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies* 74(3):705–731.
- Ganuza, J.-J., and Penalva, J. S. 2010. Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica* 78(3):1007–1030.
- Hajaj, C., and Sarne, D. 2014. Strategic information platforms: Selective disclosure and the price of free. In *Proceedings of the ACM Conference on Economics and Computation*, 839–856.
- Hajaj, C.; Hazon, N.; and Sarne, D. 2015. Improving comparison shopping agents competence through selective price disclosure. *Electronic Commerce Research and Applications* 14(6):563–581.
- Hajaj, C.; Hazon, N.; and Sarne, D. 2016. Enhancing comparison shopping agents through ordering and gradual information disclosure. *Autonomous Agents and Multi-Agent Systems* 1–19.
- Horvitz, E.; Kadie, C.; Paek, T.; and Hovel, D. 2003. Models of attention in computing and communication: from principles to applications. *Communications of the ACM* 46(3):52–59.
- Hui, B., and Boutilier, C. 2006. Who’s asking for help?: a bayesian approach to intelligent assistance. In *Proceedings of IUI*, 186–193.
- Kahneman, D. 2000. A psychological point of view: Violations of rational rules as a diagnostic of mental processes (commentary on stanovich and west). *Behavioral and Brain Sciences* 23:681–683.
- Kamar, E.; Gal, Y. K.; and Grosz, B. J. 2013. Modeling information exchange opportunities for effective human–computer teamwork. *Artificial Intelligence* 195:528–550.
- Kapoor, A.; Horvitz, E.; and Basu, S. 2007. Selective supervision: Guiding supervised learning with decision-theoretic active learning. In *IJCAI*, volume 7, 877–882.
- Keren, G. 1987. Facing uncertainty in the game of bridge: A calibration study. *Organizational Behavior and Human Decision Processes* 39(1):98–114.
- Klos, A.; Weber, E. U.; and Weber, M. 2005. Investment decisions and time horizon: Risk perception and risk behavior in repeated gambles. *Management Science* 51(12):1777–1790.
- Kraus, S. 2015. Human-agent decision-making: Combining theory and practice. In *TARK-2015*.
- Krause, A., and Guestrin, C. 2009. Optimal value of information in graphical models. *Journal of Artificial Intelligence Research* 35:557–591.
- Lai, K. P.; Chong, S. C.; Ismail, H. B.; and Tong, D. Y. K. 2014. An explorative study of shopper-based salient e-servicescape attributes: A Means-End Chain approach. *International Journal of Information Management* 34(4):517–532.
- Levy, P., and Sarne, D. 2016. Intelligent advice provisioning for repeated interaction. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 842–849.
- Liao, W., and Ji, Q. 2008. Efficient non-myopic value-of-information computation for influence diagrams. *International Journal of Approximate Reasoning* 49(2):436–450.
- Moscarini, G., and Smith, L. 2003. The optimal level of experimentation. *Econometrica* 69(6):1629–1644.
- Paolacci, G.; Chandler, J.; and Ipeirotis, P. G. 2010. Running experiments on amazon mechanical turk. *Judgment and Decision making* 5(5):411–419.
- Peled, N.; Kraus, S.; et al. 2015. A study of computational and human strategies in revelation games. *Autonomous Agents and Multi-Agent Systems* 29(1):73–97.
- Rabin, M. 1998. Psychology and economics. *Journal of Economic Literature* 11–46.
- Rysman, M. 2009. The economics of two-sided markets. *The Journal of Economic Perspectives* 125–143.
- Sarne, D.; Alkoby, S.; and David, E. 2014a. On the choice of obtaining and disclosing the common value in auctions. *Artificial Intelligence* 215:24–54.
- Sarne, D.; Alkoby, S.; and David, E. 2014b. On the choice of obtaining and disclosing the common value in auctions. *Artificial Intelligence* 215:24–54.
- Sheena, I. 2010. *The Art of Choosing*. Twelve.
- Simon, H. A. 1972. Theories of bounded rationality. *Decision and organization: A volume in honor of Jacob Marschak* 161–176.
- Thaler, R. H., and Sunstein, C. R. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
- Tolpin, D., and Shimony, S. E. 2012. Rational value of information estimation for measurement selection. *Intelligent Decision Technologies* 6(4):297–304.
- Yakout, M.; Elmagarmid, A. K.; Neville, J.; Ouzzani, M.; and Ilyas, I. F. 2011. Guided data repair. *Proceedings of the VLDB Endowment* 4(5):279–289.
- Zhang, Y. 2010. Multi-task active learning with output constraints. In *Proceeding of AAAI*.