On the Profitability of Incompetence

Eugen Staab and Martin Caminada

University of Luxembourg,
Faculty of Science, Technology and Communication,
L–1359 Luxembourg
{eugen.staab,martin.caminada}@uni.lu

Abstract. The *exchange of information* is in many multi-agent systems the essential form of interaction. For this reason, it is crucial to keep agents from providing unreliable information. However, agents that provide information have to balance between being highly competent, in order to achieve a good reputation as information provider, and staying incompetent, in order to minimize the costs of information acquisition. In this paper, we use a multi-agent simulation to identify conditions under which it is profitable for agents either to make an investment to become competent, or to economize and stay incompetent. We focus on the case where the quality of the acquired information cannot objectively be assessed in any immediate way and where hence the information end users have to rely on secondary methods for assessing the quality of the information itself, as well as the trustworthiness of those who provide it.

Keywords: Social Epistemology, Dishonesty, Formal Argumentation, Reputation Systems, Incentive Compatibility.

1 Introduction

When purchasing information, one wants to be sure of the quality of the information in question. However, if one is not an expert oneself in the relevant domain, assessing the quality of information can be difficult. For the sellers of information (which we will simply refer to as "the consultants") this provides an incentive for dishonesty. After all, gaining real expertise costs significant efforts as well as time and money. If the consumer of information (which we will refer to as "the client") has difficulties assessing the quality of the provided information, then why not pretend to have a higher level of expertise than one actually has? As long as the chance that the client detects this dishonesty is low, and so the reputation will most probably not be damaged, a consultant can charge the same price for his advice, yet spend less resources on keeping up-to-date regarding the state of the art.

The issue of low quality information has been studied in [1, 2]. What is new, however, is that we have now developed a model and a software simulator thereof that is able to compute the profit for the consultants of either a strategy of hard work or a strategy of taking it easy when it comes to staying up to date with the state of the art. In particular, we are able to provide some insight on which strategy yields the most profitable results under which circumstances.

2 Argumentation and Informedness

The aim of this section is to formalize the concept of informedness by means of formal argumentation. This establishes the background theory for the remaining practical part of the paper.

In standard epistemic logic (S5), informedness is basically a binary phenomenon. One either has knowledge about a proposition p or one does not. It is, however, also possible to provide a more subtle account of the extent to which one is informed about the validity of proposition p. Suppose Alex thinks that Hortis Bank is on the brink of bankruptcy because it has massively invested in mortgage backed securities. Also Bob thinks that Hortis is on the brink of bankruptcy because of the mortgage backed securities. Bob has also read an interview in which the finance minister promises that the state will support Hortis if needed. However, Bob also knows that the liabilities of Hortis are so big that not even the state will be able to provide significant help to avert bankruptcy. From the perspective of formal argumentation [3], Bob has three arguments at his disposal.

A**:** Hortis Bank is on the brink of bankruptcy, because of the mortgage backed securities.
B**:** The state will save Hortis, because the finance minister promised so.
C**:** Not even the state has the financial means to save Hortis.

Here, argument B attacks A, and argument C attacks B (see eq. (1)). In most approaches to formal argumentation, arguments A and C would be accepted and argument B would be rejected.

$$A \longleftarrow B \longleftarrow C \tag{1}$$

Assume that Alex has only argument A at his disposal. Then it seems reasonable to regard Bob as more informed with respect to proposition p ("Hortis Bank is on the brink of bankruptcy") since he has a better knowledge of the facts relevant for this proposition and is also in a better position to defend it in the face of criticism.

The most feasible way to determine whether someone is informed on some given issue is to evaluate whether he is up to date with the relevant arguments and is able to defend his position in the face of criticism. One can say that agent X is more informed than agent Y if it has at its disposal a larger set of relevant arguments.

We will now provide a more formal account of how the concept of informedness could be described using formal argumentation. An *argumentation framework* [3] is a pair (Ar, att) where Ar is a set of arguments and att is a binary relation on Ar. An argumentation framework can be represented as a directed graph. For instance, the argumentation framework $(\{A, B, C\}, \{(C, B), (B, A)\})$ is represented in eq. (1).

Arguments can be seen as defeasible derivations of a particular statement. These defeasible derivations can then be attacked by statements of other defeasible derivations, hence the attack relationship. Given an argumentation framework, an interesting question is what is the set (or sets) of arguments that can collectively be accepted. Although this question has traditionally been studied in terms of the various fixpoints of the characteristic function [3], it is equally well possible to use the approach of argument labelings [4, 5]. The idea is that each argument gets exactly one label (accepted, rejected, or abstained), such that the result satisfies the following constraints.

1. If an argument is labeled accepted then all arguments that attack it must be labeled rejected.
2. If an argument is labeled rejected then there must be at least one argument that attacks it and is labeled accepted.
3. If an argument is labeled abstained then it must not be the case that all arguments that attack it are labeled rejected, and it must not be the case that there is an argument that attacks it and is labeled accepted.

A labeling is called complete iff it satisfies each of the above three constraints. As an example, the argumentation framework of eq. (1) has exactly one complete labeling, in which A and C are labeled accepted and B is labeled rejected. In general, an argumentation framework has one or more complete labelings. Furthermore, the arguments labeled accepted in a complete labeling form a complete extension in the sense of [3]. Other standard argumentation concepts, like preferred, grounded and stable extensions can also be expressed in terms of labelings [4, 5]. Algorithms and proof procedures can be found in [6–11].

In essence, one can see a complete labeling as a reasonable position one can take in the presence of the imperfect and conflicting information expressed in the argumentation framework [12, 13]. An interesting question is whether an argument *can* be accepted (that is, whether the argument is labeled accepted in at least one complete labeling) and whether an argument *has to be* accepted (that is, whether the argument is labeled accepted in each complete labeling). These two questions can be answered using formal discussion games [6–8, 11]. For instance, in the argumentation framework of eq. (1), a possible discussion would go as follows.

Proponent: Argument A has to be accepted.
Opponent: But perhaps A's attacker B does not have to be rejected.
Proponent: B has to be rejected because B's attacker C has to be accepted.

The precise rules which such discussions have to follow are described in [6–9, 11]. We say that argument A can be *defended* iff the proponent has a winning strategy for A. We say that argument A can be *denied* iff the opponent has a winning strategy against A.

If informedness is defined as justified belief, and justified is being interpreted as defensible in a rational discussion, then formal discussion games can serve as a way to examine whether an agent is informed with respect to proposition p, even in cases where one cannot directly determine the truth or falsity of p in the objective world. An agent is informed on p iff it has an argument for p that it is able to defend in the face of criticism.

The thus described approach also allows for the distinction of various grades of informedness. That is, an agent X can be perceived to be at least as informed as agent Y w.r.t. argument A iff either X and Y originally disagreed on the status of A but combining their information the position of X is confirmed, or X and Y originally agreed on the status of A and in every case where Y is able to maintain its position in the presence of criticism from agent Z, X is also able to maintain its position in the presence of the same criticism.

When $AF_1 = (Ar_1, att_1)$ and $AF_2 = (Ar_2, att_2)$ are argumentation frameworks, we write $AF_1 \sqcup AF_2$ as a shorthand for $(Ar_1 \cup Ar_2, att_1 \cup att_2)$. Formally, agent X is at least as informed with respect to argument A as agent Y iff:

1. A can be defended using AF_X (that is, if X assumes the role of the proponent of A then it has a winning strategy using the argumentation framework of X), A can be denied using AF_Y (that is, if Y assumes the role of the opponent than it has a winning strategy using the argumentation framework of Y), but A can be defended using $AF_X \sqcup AF_Y$, or
2. A can be denied using AF_X, A can be defended using AF_Y, but A can be denied $AF_X \sqcup AF_Y$, or
3. A can be defended using AF_X and can be defended using AF_Y, and for each AF_Z such that A can be defended using $AF_Y \sqcup AF_Z$ it holds that A can also be defended using $AF_X \sqcup AF_Z$, or
4. A can be denied using AF_X and can be denied using AF_Y, and for each AF_Z such that A can be denied using $AF_Y \sqcup AF_Z$ it holds that A can be denied using $AF_X \sqcup AF_Z$.

In the example mentioned earlier (eq. (1)) Alex has access only to argument A, and Bob has access to arguments A, B and C. Suppose a third person (Charles) has access only to arguments A and B. Then we say that Bob is more informed than Alex w.r.t. argument A because Bob can maintain his position on A (accepted) while facing criticism from Charles, where Alex cannot. A more controversial consequence is that Charles is also more informed than Alex w.r.t. argument A, even though from the global perspective, Charles has the "wrong" position on argument A (rejected instead of accepted). This is compensated by the fact that Bob, in his turn, is more informed than Charles w.r.t. argument A. As an analogy, it would be fair to consider Newton as more informed than his predecessors, even though his work has later been attacked by more advanced theories.

It can be interesting to compare the thus defined notion of argumentation-based informedness with the notion of knowledge as modeled by traditional (S5) modal logic. Knowledge, from a conceptual point of view, is often defined as "justified true belief". When using S5 and S4 based modalities, the notion of knowledge is usually simplified as "true belief", whereas in our argumentation approach, we take the other way and define informedness as "justified belief". The difference between the modal logic approach and the argumentation approach is an important one, since it has consequences for the domains where these approaches are applicable. As an example, consider an expert on climate change who predicts a global temperature increase of 2^o C by the year 2050. Whether or not this claim is true or not cannot immediately be assessed in any objective way. However, what can be assessed is whether the backing of this claim can stand a critical assessment using the information that is currently available. That is, is the expert able to defend his position against possible counterarguments? Similar observations can be made not only with respect to climate change, but also with respect to issues like the world's energy resources, or the viability of the long-term investment strategy of a pension fund. The reputation of the experts who work in these fields cannot be purely determined in terms of feedback from the objective world, since in many occasions this feedback will only reveal itself at the end of one's professional life. In many cases one cannot determine whether a statement is *true*; one can only determine whether it is *well-informed*.

3 Model

We consider a client/consultant-scenario, that is, a scenario where consultants advise their clients on a certain issue. We model the knowledge on which the consultants advise their clients by a chain of arguments:

$$A_1 \longleftarrow A_2 \longleftarrow \ldots \longleftarrow A_{N_{\mathrm{arg}}} \tag{2}$$

Here, any argument A_i (for $1 < i \leq N_{\mathrm{arg}}$) defeats its predecessor argument A_{i-1}. As a consequence, if N_{arg} is even, then all arguments A_i with even indices are accepted, and all arguments with odd indices are rejected. For odd N_{arg}, it is the other way around.[1]

At the beginning of a simulation, only argument A_1 is known to the consultants and only this argument is known in the whole society, i.e., it represents the "state of the art". To model the discovery/emergence of new information (e.g., through research), we make a certain number of new arguments available to the consultants in each round. This represents the evolution of the state of the art. The number of new arguments per round will be fixed for a simulation and is denoted by ΔN_{arg}. The simulation is finished when all N_{arg} arguments have been made available. During simulation, the structure of the chain of arguments looks as follows ($k \leq i$ must hold):

$$\overbrace{\underbrace{A_1 \leftarrow \cdots \leftarrow A_k}_{\text{known to a certain consultant}} \leftarrow \cdots \leftarrow A_i}^{\text{"state of the art"}} \leftarrow \overbrace{A_{i+1} \leftarrow \cdots \leftarrow A_{i+\Delta N_{\mathrm{arg}}}}^{\text{added to the "state of the art" in the next round}} \leftarrow \cdots \leftarrow A_{N_{\mathrm{arg}}} \tag{3}$$

In each round, consultants can decide how many new arguments they want to procure. We assume that the consultants extend their already known chain of arguments with new arguments always in a seamless manner, i.e., without gaps. This assumption was made in order to be in line with argument games (such as described in [6–8]) where each uttered argument is a reaction to a previously uttered argument, thus satisfying the property of *relevance* [14].

[1] Although it would have been possible, and to some extent even more natural, to use a tree-shaped argumentation framework instead of just a linear one, we do not expect our current simplification to significantly affect the outcome of the simulator.

3.1 Expenses, Turnover and Profit

For the sake of simplicity, we model the cost of an argument by some constant c_{arg}. This means that to get for instance the knowledge about argument A_{10}, a consultant has an overall expense of $10 \cdot c_{\text{arg}}$ (recall that arguments can only be procured in a row). We write n_{arg} to denote the total number of arguments acquired by a specific consultant (where $n_{\text{arg}} \leq N_{\text{arg}}$). Then, the expenses E of a consultant can be computed as:

$$E = n_{\text{arg}} \cdot c_{\text{arg}} \tag{4}$$

The turnover of a consultant is defined as the sum of the money that the consultant has been paid. Of course, the consultant is paid only for those consultations where he actually is better informed than the client; we call these "successful consultations". Let S be the multiset that contains all amounts that have been paid to a certain consultant. This consultant's turnover T is defined as:

$$T = \sum_{p \in S} p \tag{5}$$

The profit P of a consultant is defined as the difference between his turnover and his expenses:

$$P = T - E \tag{6}$$

3.2 Consultancy Strategies

Consultants generally want to provide as little information as necessary, because this way they can give more consultations. At the same time, consultants want to give advice that makes them appear knowledgeable – in order to increase their reputation. Therefore, in our model, a consultant advises a client always with the argument that has the lowest index above the client's knowledge and that is compliant with the consultant's latest known argument, i.e., that has the same parity. In other words, provided that a consultant knows enough arguments, he provides a client with *two* arguments, if the latest argument known to the client is of the same parity as the latest argument known to the consultant, and with *one* argument otherwise. These arguments become known to the client.

We consider two strategies for how consultants can increase their knowledge:

Well-informed strategy *(WELL)*: A consultant procures arguments as soon as these become available, so as to be always up-to-date with the aim to achieve a good reputation.

Ill-informed *(ILL)*: A consultant procures arguments only as to appear knowledgeable to the clients. More precisely, only upon encountering a client who is as informed as the consultant (before or after the consultation), or even better informed, the consultant procures a number of new arguments, which we set to 2. Although this strategy could be made much more sophisticated, we show that under certain conditions it outperforms the *WELL*-strategy already in this form.

Consultants that follow the *WELL*-strategy are always as competent as possible, whereas consultants that follow the *ILL*-strategy become increasingly incompetent with increasing ΔN_{arg}. The *ILL*-strategy allows consultants to offer their advice at a lower price, because they have to invest less in new information. However, this comes at the cost of risking a decrease in reputation, because clients do not want to be advised by a consultant who is not better informed than they are.

3.3 Selection of Consultants

In our model, clients rate consultants according to two criteria: the *price* demanded by the consultants, and their *reputation*.

Price: Clients prefer relatively cheap consultants. The price is agreed upon by client and consultant before an interaction takes place.

Reputation: A client wants to get advice from consultants with a good reputation. In this context, reputation reflects the characteristics of the consultant that cannot be agreed upon beforehand, because they can generally not be checked after an interaction. For instance, in our scenario, clients are generally unable to check provided information for correctness.

We denote a consultant i's current reputation by r_i and represent his price for the upcoming round in form of "cheapness", denoted by c_i. The details on how the reputation and cheapness are computed in our model are given later. For now, it suffices to know that both values are in the interval $(0, 1]$. A high cheapness and a high reputation make a consultant attractive. A parameter $\alpha \in [0, 1]$ defines which of the two criteria the clients think is more important. The "attractiveness" a_i of consultant i is defined as (and is recomputed each round):

$$a_i = \alpha \cdot c_i + (1 - \alpha) \cdot r_i \tag{7}$$

A high α favors cheaper consultants, while a low α favors more reputable consultants. In each round, each client selects a new consultant. Attractiveness values are first centered around a mean of 0.5 (to weaken the impact of extreme outliers), and then normalized to $[0, 1]$, giving a'_i. Finally, a client selects consultant i with the following probability:[2]

$$P_i = \frac{a'_i}{\sum_j a'_j} \tag{8}$$

If a client meets a consultant who is not better informed than he is, the client repeats the selection procedure.

[2] In the implementation, we reserve for each consultant i a disjoint interval with length P_i, and generate for each client a uniform random number that selects his consultant: by the interval it falls on (note that a consultant can be selected by several clients).

Price Computation Let δ be the *profit margin* of a consultant, with $\delta \in [0, \infty)$, where $\delta = 0.5$ represents for instance a profit margin of 50%. Using a certain profit margin δ, a consultant i computes his current price p_i as follows:

$$p_i = (1 + \delta) \frac{E}{|S|} \tag{9}$$

Here, $\frac{E}{|S|}$ is a heuristic to provide cost recovery, where E models the expenses (see eq. (4)), and $|S|$ is the number of successful consultations so far (see eq. (5)). Still, no client would choose a consultant that is more expensive than the acquisition of the information itself. Hence, we limit the price to the cost of *one* argument (see also Sect. 3.2). We map each price to the interval $(0, 1]$ and transform it into cheapness c_i as follows:

$$c_i = \frac{\min_j(p_j)}{p_i} \tag{10}$$

In this way, the cheapest consultant has maximal cheapness 1, and the ratios between the prices are preserved, as can easily be shown:

$$\forall i, j : \frac{c_i}{c_j} = \frac{\min_k(p_k)}{p_i} \cdot \frac{p_j}{\min_k(p_k)} = \frac{p_j}{p_i} \tag{11}$$

Reputation Computation In our model, clients use a *reputation system* [15] to share their experiences with consultants. This allows clients to better estimate the trustworthiness of the consultants and thus to better select their future consultants. We assume "perfect" conditions for the reputation system, because this will make it harder for the consultants with the *ILL*-strategy to hold their ground. These perfect conditions consist of:

– *honest reporting* of the clients, i.e., clients do not bias their experience,
– all clients have the same idea of how to fuse the experiences with consultants, and so a *global reputation score* can be computed, and
– *total information sharing*, i.e., every client shares *all* his experiences with every other client.

To minimize the impact of specifics of the reputation system on our results, we try to keep it as simple as possible. We propose a system that measures the reputation of a consultant based on the number of bad and good experiences with that consultant. Because clients cannot verify the arguments, they have a bad experience with a consultant only if the consultant is not better informed than they are. Such an interaction is evidence for a consultant following the *ILL*-strategy; in rare cases, this interaction can also be misleading evidence, namely in the case where the consultant is actually following the *WELL*-strategy and the client's knowledge is state of the art. How often the evidence is misleading depends on how fast new information becomes available (ΔN_{arg}); for $\Delta N_{\text{arg}} \geq 3$ for instance, the consultants that follow the *WELL*-strategy are always ahead of the clients, and so a bad experience implies an encounter with a consultant following the *ILL*-strategy. The clients share their experience and maintain for each consultant i a global counter $\mathcal{G}_i$ of good experiences, and a global counter $\mathcal{B}_i$ of bad experiences. Then a reputation score is computed as follows (we follow the trust value computation from [16]):

$$r'_i = \frac{\mathcal{G}_i + 1}{\mathcal{G}_i + \mathcal{B}_i + 2} \tag{12}$$

It follows that at the point where no experience with a consultant has been made yet ($\mathcal{G}_i = \mathcal{B}_i = 0$), his reputation is 0.5. To make reputation comparable to cheapness, we map it to $(0, 1]$ as follows:

$$r_i = \frac{r'_i}{\max_j(r'_j)} \tag{13}$$

As for cheapness, the most reputable consultant has reputation 1, and ratios between reputation scores are preserved (proof analogously to eq. (11)).

4 Simulations

We have implemented a simulator for our model. The aim of this simulator is to reveal the impact of the different model parameters on the profit of the two consultancy strategies. In other words, we want to identify the parameter settings for which the *ILL*-strategy is more profitable than the *WELL*-strategy.

4.1 Experiments

Each experiment was repeated 2^{10} times. Mean profits and corresponding standard deviations were computed, separately for the two consultancy strategies. To account for the fact that consultancy makes only sense when a consultant can advise several clients, we chose a much higher number of clients (2^{10}) than consultants (2^7). The sets of clients and consultants are fixed. At the outset, all consultants procure 2 initial arguments.[3] We varied the following parameters:

- the number of arguments becoming available each round ($\Delta N_{\text{arg}} \in \{2, 4, 6, 8\}$),
- the fraction of consultants that use the *ILL*-strategy ($f_{ILL} \in \{0.1, 0.5, 0.9\}$), where the remaining consultants use the *WELL*-strategy,
- the profit margin ($\delta \in \{0.1, 0.5\}$), and
- the factor α that regulates the importance of the consultants' price and reputation for the clients ($\alpha \in [0, 1]$).

[3] We also run the experiments with all consultants procuring 4 initial arguments. As a result, the profit of the *ILL*-consultants increased in all cases. Because of the limited space, these results are not shown here.

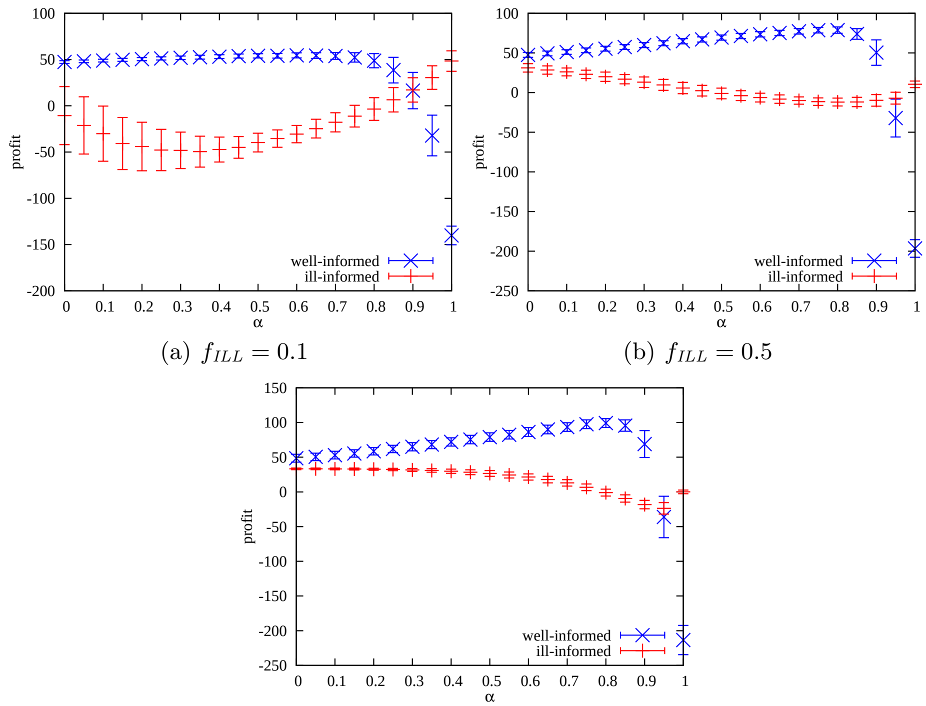

(a) $f_{ILL} = 0.1$

(b) $f_{ILL} = 0.5$

(c) $f_{ILL} = 0.9$

Fig. 1. Impact of fraction of *ILL*-consultants ($\Delta N_{\text{arg}} = 2$ and $\delta = 0.1$).

4.2 Results

In the following, we analyze the results of selected experiments.[4] In all figures, the x-axis gives α, which defines how clients chose their consultants, i.e., for $\alpha = 0$, clients select consultants solely based on their reputation, and attach more importance to the price for increasing α; for $\alpha = 1$, clients only look at the price of the consultants (see also Sect. 3.3). The y-axis gives the profit of the consultants for both consultancy strategies.

Impact of Fraction of *ILL*-Consultants At the outset, we look at the impact of f_{ILL} on the consultants' profit. To this end, we first fix the parameters $\Delta N_{\text{arg}} = 2$ and $\delta = 0.1$. The results that are shown in Fig. 1 reveal that for small α, the profit of the two types of consultants converge for increasing f_{ILL}, whereas for large α (with a center roughly around 0.8) they develop in different directions. The profit of *ILL*-consultants is in certain areas very low — and even

[4] Another set of results was presented in [17]. However, the current model is more reasonable in fundamental points like price and reputation computation, and thus the results of the old model are not considered here.

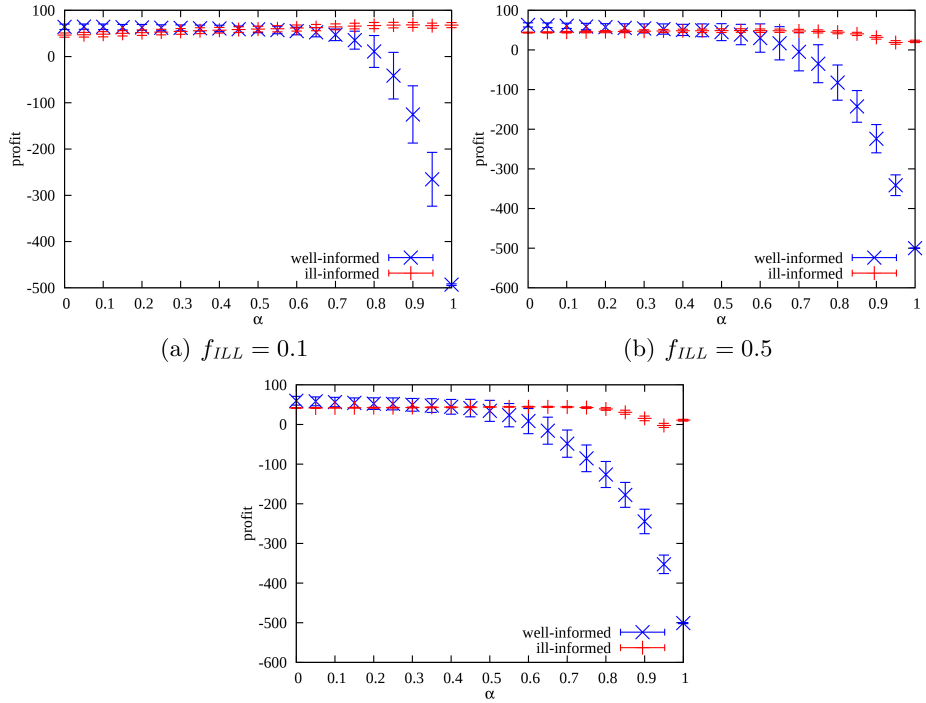

(a) $f_{ILL} = 0.1$ (b) $f_{ILL} = 0.5$

(c) $f_{ILL} = 0.9$

Fig. 2. Impact of fraction of *ILL*-consultants ($\Delta N_{\text{arg}} = 4$ and $\delta = 0.1$).

negative. By looking at the data, we found that the reason for this is their bad reputation (due to a high rate of unsuccessful consultations, for which they are not paid), and their price which is not considerably better in this scenario. For increasing f_{ILL}, these "negative areas" seem to shift to the right: we found in the data that the average attractiveness of *ILL*-consultants is increasing for low α, while for the *WELL*-consultants it is increasing for higher α (up to a certain point of α). This an effect of a complex interplay of price, reputation and selection, and we currently have no exact explanation for this. For very high α, one can see a drop in the profit of *WELL*-consultants, because the price is becoming decisive for selection here.

As can be seen from Fig. 2, where ΔN_{arg} is higher, an increase in f_{ILL} causes the drop of the profit of *WELL*-consultants for high α to become more intense. The reason is that now it becomes harder for *WELL*-consultants to offer competitive prices, and so, since they get more competitors for increasing f_{ILL}, it becomes harder for them to make profit, especially when clients care much about the price. For the same reason the profit of *ILL*-consultants is (slightly) decreasing for high α and increasing f_{ILL}: a higher number of *ILL*-consultants

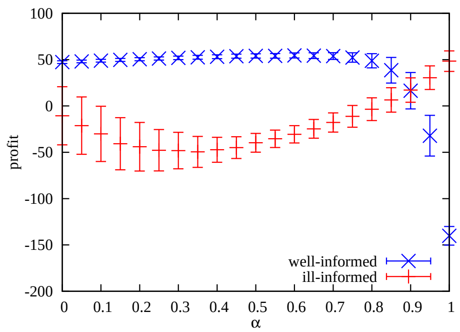

(a) $\Delta N_{\text{arg}} = 2$

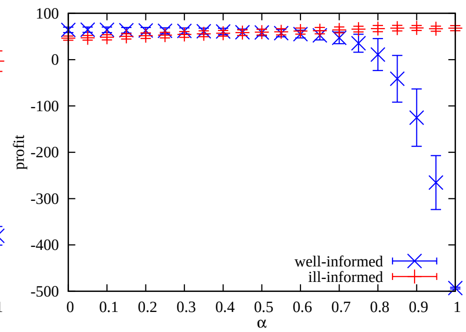

(b) $\Delta N_{\text{arg}} = 4$

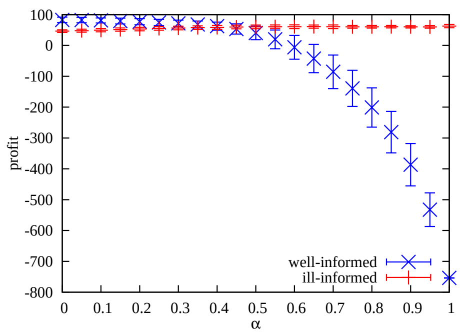

(c) $\Delta N_{\text{arg}} = 6$

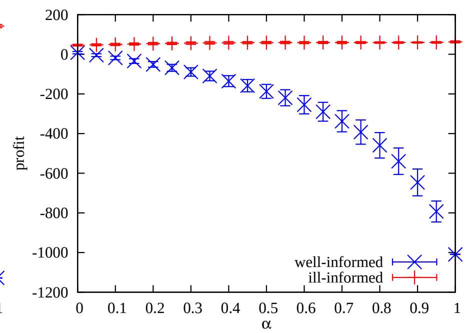

(d) $\Delta N_{\text{arg}} = 8$

Fig. 3. Impact of ΔN_{arg} ($\delta = 0.1$, $f_{ILL} = 0.1$).

has to share the profit. Still, their profit increases relatively to the profit of the *WELL*-consultants.

Note that in Figure 1(c) for $\alpha = 0.95$, both strategies make negative profit. This is due to the limitation of the price to the price of one argument (see Sect. 3.3), which is in this case not sufficient to provide cost recovery.

Impact of $\boldsymbol{\Delta N_{\text{arg}}}$ We now look at the impact of a higher speed of arguments becoming available (ΔN_{arg}). Figure 3 shows results for varying ΔN_{arg} and fixed δ and f_{ILL}. It can be observed that the profit of the *ILL*-strategy increases in comparison to the *WELL*-strategy for increasing ΔN_{arg}. The reason for this is that for a higher ΔN_{arg}, the *WELL*-consultants have to invest more in the arguments to keep up with the state of the art, and thus are more expensive. Being selected less often, they have to ask for higher prices to compensate their loss (they are continually procuring arguments). For $\Delta N_{\text{arg}} = 8$, that even goes so far as to make it in general unprofitable to follow the *WELL*-strategy, independent of the clients' preferences α. At the same time, the *ILL*-consultants have more successful consultations, and thus are paid more. This is because for in-

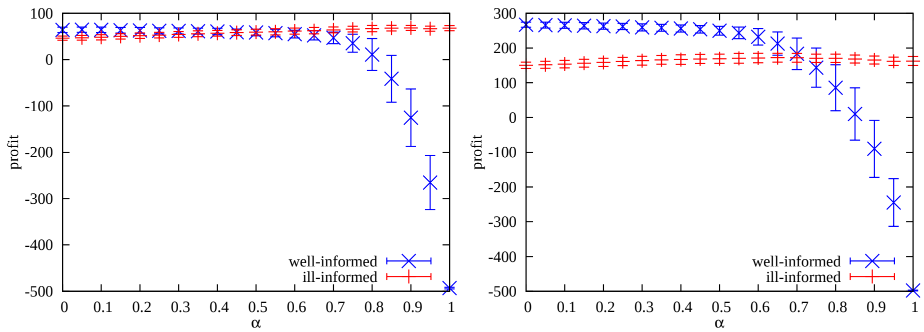

(a) $\delta = 0.1$ (b) $\delta = 0.5$

Fig. 4. Impact of profit margin ($\Delta N_{\text{arg}} = 4$, $f_{ILL} = 0.1$).

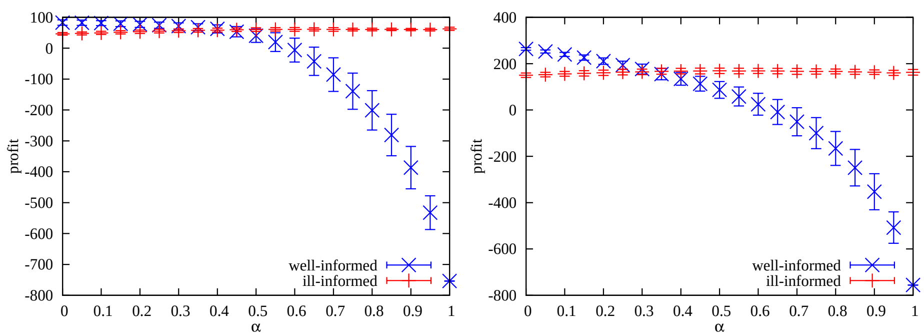

(a) $\delta = 0.1$ (b) $\delta = 0.5$

Fig. 5. Impact of profit margin ($\Delta N_{\text{arg}} = 4$, $f_{ILL} = 0.5$).

creasing ΔN_{arg}, *ILL*-consultants are chosen more regularly by the clients, and so are better informed about the informedness of the clients.

Impact of Profit Margin Up to now, we have looked at a profit margin of $\delta = 0.1$. Figures 4, 5 and 6 show what happens when δ is increased; the left figures show $\delta = 0.1$, the right figures show $\delta = 0.5$. In general, it can be seen that a higher profit margin increases the maximal profit (e.g., look at the scale of the y-axis). Apart from that, in Figures 4 and 5, the profit of the *WELL*-strategy is increased for low α, but not affected much for high α. The point where the increase ceases to take place, seems to be where the two profit curves intersect. This holds also for Fig. 6 where the two profit curves move away from each other (there is no intersection). It is also confirmed by the other results not shown in this paper. This leads us to the conjecture that the profit margin amplifies the difference between the profit of the two types of consultants.

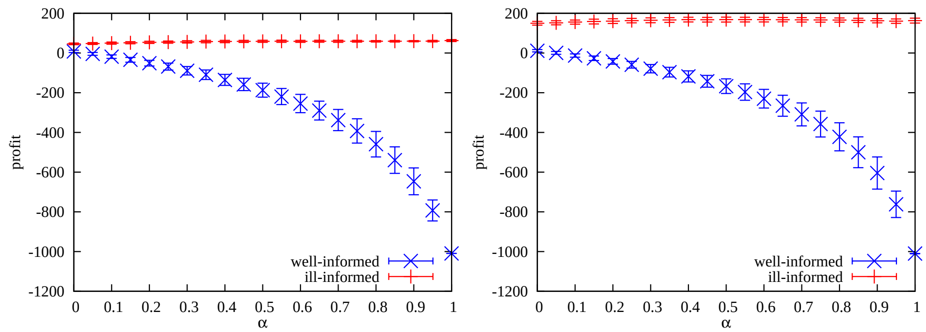

(a) $\delta = 0.1$ (b) $\delta = 0.5$

Fig. 6. Impact of profit margin ($\Delta N_{\mathrm{arg}} = 8$, $f_{ILL} = 0.1$).

5 Conclusions & Perspectives

In this paper, we have compared the profit that consultants yield when following two different strategies: staying as competent as possible, and staying preferably incompetent. In our model, we found that there are scenarios where it is more profitable to stay incompetent. In particular, this is increasingly the case when:

- the speed with which the state of the art changes (ΔN_{arg}) is high,
- clients prefer cheap to reputable consultants (there are exceptions when both ΔN_{arg} is small and f_{ILL} is high), or
- the fraction of incompetent consultants is high (and ΔN_{arg} is not too low).

The impact of the profit margin seems to be more complex. However, it appears to act as an amplifier in that it increases the difference between the profit of the two consultancy strategies.

Our finding that considerations of reputation are sometimes not sufficient to counterbalance incentives for providing low quality information has also been observed in the field of economics, where the role of credit rating agencies has received significant criticism, before as well as after the recent credit crisis. Mathis et al. claim that considerations of reputation might not be sufficient to dissuade rating agencies from giving a too positive rating to certain structured products [18]. Other criticism regarding the low added value of rating agencies has been provided by [19–21]. Despite the highly questionable performance of the credit rating agencies, as witnessed in the recent credit crisis, it is striking to see that the ratings industry has continued to operate, while still being relied on both by regulators and by the public [22]. Although some important differences between the domain of credit ratings and the domain studied in our software simulator do exist (for instance the low number of rating agencies compared to the relatively high number of consultants in our simulator) it is striking to see that our findings on the profitability of incompetence do not fundamentally deviate from what has been observed in economics.

The reputation system implemented in the current paper is relatively simple and straightforward. An alternative would be to implement a more advanced approach, like [23, 24], which essentially uses a trust-net for weighting the other agents' opinions. Where in [23, 24] reporting to the reputation system is assumed to be honest, ReGret [25, 26] goes one step further and considers the possibility of dishonest reporting. In our current work, however, the problem studied is not so much dishonest reporting (we assume all client agents to be honest) but fundamental limitations regarding the extent to which a client is able to evaluate the quality of the advise it has purchased. That is, we show that even when all reporting to the reputation system is honest, it can still be the case that the *ILL*-consultants achieve the best economic performance. Nevertheless, the reputation system could be made more clever. For instance, clients could look retrospectively at what their consultants advised them on: is this *older advice* conflicting with what a client *currently believes*? Depending on how much the client believes the real world has changed in between, he can retroactively reduce the reputation of the respective consultants – which would not be justified though if his current believes are incorrect.

There are many ways of how to extend our current model, and to make it more realistic. First of all, the price computation could be extended with ideas from the field of economics: instead of caring just about cost recovery, consultants could proactively reduce their price in order to attract more clients; this would involve models for market analysis. Also, the possibility of bankruptcy could be considered. Furthermore, in our model, clients select their consultants autonomously, but all have the same preference regarding price and reputation. This can be extended by defining a probability density function (PDF) over the clients' preferences, and choosing the preference of a client from this distribution. Domain specific knowledge should be used to define a meaningful PDF. Analogously, the number of arguments getting available each round to the consultant (ΔN_{arg}) could be described by a PDF. Also, this work did not address the issue of dynamics, e.g., clients that adapt their preferences over time. Finally, more complex argumentation frameworks could be explored, possibly involving trees instead of the relatively simple linear structure treated in the current paper.

Thanks

We would like to thank Yining Wu (Interdisciplinary Lab for Intelligent and Adaptive Systems) and Tibor Neugebauer (Luxembourg School of Finance) for their useful comments.

References

1. Frankfurt, H.G.: On Bullshit. Princeton University Press (2005)
2. Caminada, M.: Truth, lies and bullshit; distinguishing classes of dishonesty. In: Social Simulation Workshop at the International Joint Conference on Artificial Intelligence (SSIJCAI). (2009)

3. Dung, P.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence **77** (1995) 321–357
4. Caminada, M.: On the issue of reinstatement in argumentation. In Fischer, M., van der Hoek, W., Konev, B., Lisitsa, A., eds.: Logics in Artificial Intelligence; 10th European Conference, JELIA 2006, Springer (2006) 111–123 LNAI 4160.
5. Caminada, M., Gabbay, D.: A logical account of formal argumentation. Studia Logica **93**(2-3) (2009) 109–145 Special issue: new ideas in argumentation theory.
6. Prakken, H., Sartor, G.: Argument-based extended logic programming with defeasible priorities. Journal of Applied Non-Classical Logics **7** (1997) 25–75
7. Vreeswijk, G., Prakken, H.: Credulous and sceptical argument games for preferred semantics. In: Proceedings of the 7th European Workshop on Logic for Artificial Intelligence (JELIA-00). Number 1919 in Springer Lecture Notes in AI, Berlin, Springer Verlag (2000) 239–253
8. Caminada, M.: For the sake of the Argument. Explorations into argument-based reasoning. Doctoral dissertation Free University Amsterdam (2004)
9. Modgil, S., Caminada, M.: Proof theories and algorithms for abstract argumentation frameworks. In Rahwan, I., Simari, G., eds.: Argumentation in Artificial Intelligence. Springer (2009) 105–129
10. Caminada, M.: An algorithm for computing semi-stable semantics. In: Proceedings of the 9th European Conference on Symbolic and Quantitalive Approaches to Reasoning with Uncertainty (ECSQARU 2007). Number 4724 in Springer Lecture Notes in AI, Berlin, Springer Verlag (2007) 222–234
11. Caminada, M., Wu, Y.: An argument game of stable semantics. Logic Journal of IGPL **17**(1) (2009) 77–90
12. Caminada, M., Pigozzi, G.: On judgment aggregation in abstract argumentation. Journal of Autonomous Agents and Multi-Agent Systems (2009)
13. Rahwan, I., Tohmé, F.: Collective argument evaluation as judgement aggregation. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010). (2010)
14. Caminada, M.: Semi-stable semantics. In Dunne, P., Bench-Capon, T., eds.: Computational Models of Argument; Proceedings of COMMA 2006, IOS Press (2006) 121–130
15. Resnick, P., Kuwabara, K., Zeckhauser, R., Friedman, E.: Reputation systems. Commun. ACM **43**(12) (2000) 45–48
16. Staab, E., Fusenig, V., Engel, T.: Towards trust-based acquisition of unverifiable information. In: Cooperative Information Agents XII. Volume 5180 of LNAI., Springer-Verlag (2008) 41–54
17. Staab, E., Caminada, M.: Assessing the impact of informedness on a consultant's profit. Technical report, University of Luxembourg (2009) Online: `http://arxiv.org/abs/0909.0901`.
18. Mathis, J., McAndrews, J., Rochet, J.: Rating the raters: are reputation concerns powerful enough to discipline rating agencies? Journal of Monetory Economics **56** (2009) 657–674
19. G. Ferru, L.L., Stiglitz, J.: The procyclic role of rating agencies: evidence from the east asian crisis. Economic Notes **28**(3) (1999) 335–355
20. Sinclair, T.: The new masters of capital; American bond rating agencies and the politics of creditworthyness. Cornell University Press (2005)
21. Sy, A.: Rating the rating agencies: anticipating currency crises or debt crises? Journal of Banking & Finance **28** (2004) 2845–2867

22. Spatt, C.: Discussions of 'ratings shopping and asset complexity: A theory of rating inflation'. Journal of Monetary Economics **56** (2009) 696–699
23. Schillo, M., Funk, P., Rovatsos, M.: Using trust for detecting deceitful agents in artificial societies. Applied Artificial Intelligence **14**(8) (2000) 825–848
24. Yu, B., Singh, M.P.: A social mechanism of reputation management in electronic communities. In Klusch, M., Kerschberg, L., eds.: Cooperative Information Agents IV - The Future of Information Agents in Cyberspace. Volume 1860 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2004) 355–393
25. Sabater, J., Sierra, C.: Regret: reputation in gregarious societies. In: Agents. (2001) 194–195
26. Sabater, J., Sierra, C.: Reputation and social network analysis in multi-agent systems. In: AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems, New York, NY, USA, ACM (2002) 475–482