Representative Agents and the Contract Negotiation Cold Start Problem

Federico Cerutti, Christopher Burnett, and Nir Oren

Department of Computing Science
University of Aberdeen, UK

Abstract. Principal Agent Theory (PAT) seeks to identify the incentives and sanctions that a consumer should apply when entering into a contract with a provider in order to maximise their own utility. However, identifying *suitable* contracts — maximising utility while minimising regret — is difficult, particularly when little information is available about provider competencies. In this paper we show that a global contract can be used to govern such interactions, derived from the properties of a representative agent. After describing how such a contract can be obtained, we analyse the contract utility space and its properties. Then, we show how this contract can be used to address the cold start problem and that it significantly outperforms other approaches. Finally, we discuss how our work can be integrated with existing research into multi-agent systems.

1 Introduction

Autonomous agents are often assumed to be rational, self interested entities, interacting with others in order to maximise their own utility. When asked to fulfil a task, they will therefore do so in a way that maximises their expected utility. When acting as a service provider (e.g., in an electronic marketplace), there is thus a risk that the agent will provide a substandard service. Approaches to mitigate this risk includes the use of electronic contracts [17, 18], which specify the rewards and penalties (or more generally, *incentives*) to be imposed on interacting parties in response to successful or unsuccessful interactions [6]. Principal Agent Theory (PAT) [9, 11, 19, 16] aims to determine the optimal level of incentives — in the form of rewards and penalties — that an agent (the *principal* or *consumer*) must commit to giving others (the *providers*) in order to have the latter act in such a way so as to maximise the principal's utility. To utilise PAT an agent requires beliefs about the behaviour of the provider. However, without previous (potentially negative) experiences, such beliefs cannot be formed. This problem, of lack of experience with others in the system potentially leading to poor experiences when operating within the system, is referred to as the *cold start* problem [23]. Several approaches have been proposed for addressing the cold start problem, from minimal expectation or random assignment to the capabilities of the providers [21, 24], to active learning [14]. In particular, good results have been shown by using samples of the society [22].

Apart from the cold start problem, several other difficulties arise when using PAT. Computing incentives requires solving a highly non-linear optimisation problem. When combined with the need to select between multiple possible providers, the computational costs of creating contracts using PAT, and gathering the information needed to create such contracts, become prohibitive. When dealing with unfamiliar parties humans often resort to general principles to determine incentives, stemming from cultural, psychological or legal foundations. In this paper, we build on this intuition, suggesting that without additional information, an approximate set of incentives can be specified *for all interactions within the system*. We envision that a PAT based system would initially utilise this approximate set of incentives to generate contracts. As more information becomes available through repeated interactions, these approximations become discounted in favour of more accurate incentives to form better contracts. However, in the case of very simple computationally bound agents, our approximations could continue to be used. Our work can therefore be seen to address the cold start problem by allowing an agent to successfully interact with others in the absence of specific information about them.

Our contributions are as follows. We describe a procedure for determining suitable approximate global incentive values. Such incentives aim to be applicable to all agents in the system, and in defining them, we consider their effects on overall system utility, which we refer to as the *social utility*. We define a set of incentives, or a *contract*, as *suitable* if it is the result of a trade-off between the social utility that can be gained, and the regret of paying too much for a given good or task. Informally, the contract is based on an average individual provider computed from the profile of all agents in the system. We experimentally evaluate our contribution, showing how using the global contract significantly outperforms other techniques aimed at addressing the cold start problem.

In the next section, we provide some background on principal agent theory, following which we describe how global contracts are computed in Section 3. Section 4 discusses another set of experiments evaluating the performance of using the global contract in solving the cold start problem. Section 5 summarises our results and concludes.

2 Background and Assumptions

2.1 Preliminary Notions

Following [5], we take as given a *society* of agents $\mathcal{A} = \{x, y, \ldots\}$ and a *set of tasks* T. A *consumer* $x \in \mathcal{A}$ desires to see some *task* $\tau \in T$ accomplished and must do so by having a *provider* $y \in \mathcal{A}$ perform the task on its behalf. Given $\tau \in T$, let $\mathcal{O}_\tau = \{o_0, o_1, o_2, \ldots, o_n\}$ denote the set of possible *outcomes* for task τ, where $o_0 \equiv abs$ represents the case where the provider abstained from executing the task. $\succ_o$ induces a total strict order over $\mathcal{O}_\tau$, such that intuitively, if $o_i \succ_o o_j$, o_i is *better* than o_j. $\min_o \mathcal{O}_\tau$ represents the worst possible outcome of task τ, i.e., *complete failure*, while $\max_o \mathcal{O}_\tau$ represents *complete success*. For ease of notation, $o_i \prec_o o_j$ iff $o_j \succ_o o_i$. We assume that all agents share the same task evaluation criteria as well as the same ordering function.

In *delegating* a task to a provider, the consumer asks the provider to execute it. The delegation of a task results in the consumer and provider obtaining some utility (for the consumer, due to the execution of the desired task, and for the provider, due to payment obtained from the consumer). Given this, the utility gained by the consumer is computed by the function $U^x : \mathcal{O}_\tau \mapsto \mathbb{R}$, while the provider gains utility $V^y : \mathcal{O}_\tau \mapsto \mathbb{R}$.

The task provider has autonomy in selecting the method by which a task will be carried out. In particular, let $\mathcal{E}_\tau = \{e_0, e_1, \ldots, e_m\}$ denote the set of *effort levels* they can apply when performing τ, where $e_0 \equiv abs$ identifies the case where the provider abstains from performing the task. We define a total ordering $\succ_e$ over $\mathcal{E}_\tau$ such that if $e_i \succ_e e_j$, then e_i requires more effort (or is *higher* than e_j. Similarly as before, $e_i \prec_e e_j$ iff $e_j \succ_e e_i$.

Each effort has an associated cost determined by the function $Cost : \mathcal{A} \times \mathcal{E}_\tau \mapsto \mathbb{R}$. For ease of notation, for agent $y \in \mathcal{A}$, $Cost^y$ denotes its cost function. Now different effort levels have an impact on outcomes, which we capture through a probability distribution: $\forall o \in \mathcal{O}_\tau, \forall e \in \mathcal{E}_\tau, p^y(o \mid e)$ represents the probability that agent (provider) y will achieve the outcome o using effort e. It is assumed that $p^y(abs \mid abs) = 1$.

When delegating, the consumer devises a payment function, or *contract*, $\mathcal{C} : \mathcal{A} \times \mathcal{A} \times T \times \mathcal{O} \mapsto \mathbb{R}$. We write $C^x_{y:\tau}(o)$ for a given $o \in \mathcal{O}_\tau$ to represent the contract specifying the compensation consumer x will give to provider y given outcome o of task τ.

Therefore, the *net utility* nV^y for a provider y which achieves an outcome given a specific effort (including abstention) is:

$$nV^y(o, e) = V^y(o) + C^x_{y:\tau}(o) - Cost(e) \tag{1}$$

2.2 A Fair System

In what follows we assume a *fair system* with the following properties.

(F_1): $\forall o_i, o_j \in \mathcal{O}_\tau$, if $o_i \succ_o o_j$ then $U^x(o_i) \geq U^x(o_j)$ and $V^y(o_i) \geq V^y(o_j)$;
(F_2): $\exists o_i \in \mathcal{O}_\tau$ s.t. $V^y(abs) < V^y(o_i)$;
(F_3): $\forall e_i, e_j \in \mathcal{E}_\tau$, if $e_i \succ_e e_j$ then $Cost(e_i) \geq Cost(e_j)$;
(F_4): $\forall o_i, o_j \in \mathcal{O}_\tau$, if $o_i \succ_o o_j$ then $C^x_{y:\tau}(o_i) \geq C^x_{y:\tau}(o_j)$;
(F_5): $C^x_{y:\tau}(abs) = Cost(abs)$.

(F_1) states that the better the outcome of a task, the greater the utility that (independently) both the provider and the consumer receive; (F_2) that the utility gained from abstaining cannot be smaller than the utility gained from *any* outcome; (F_3) that the higher the effort, the higher the associated cost to the provider; (F_4) that the better the outcome, the higher the compensation to the provider according to the contract (incentive); (F_5) that the contract covers the costs associated to abstaining behaviour, but no more. We note that these constraints are not minimal.

Although in reality some of those properties might fail from being satisfied, they try to capture a minimal set of norms for a free-market society. Far from being unquestionable, we elicit them as postulates on which we base our proposal in the next sections.

2.3 Rationality Assumptions

We assume that each provider rationally decides whether or not to accept a contract, and which effort to use if it does not abstain. In particular, if the provider's expected utility is greater than the utility it would obtain abstaining, then the provider will perform the requested task. Moreover, the provider will utilise the effort on the task which will maximise its own expected utility. EV^y denotes the expected utility for a provider y in performing a task τ with a contract $C^x_{y:\tau}$, and is computed as follows.

$$EV^y = \sum_{e \in \mathcal{E}_\tau} \sum_{o \in \mathcal{O}_\tau} \big(p(o|e)\ (V^y(o) + C^x_{y:\tau}(o) - \mathit{Cost}(e))\big) \tag{2}$$

Therefore, if $EV^y \leq V^y(abs)$, it is better, or *more convenient* for the provider to abstain from accepting the task. If, instead, $EV^y > V^y(abs)$, then the effort that the provider will expend on performing the task is as follows.

$$\operatorname*{argmax}_{e \in \mathcal{E}_\tau} \sum_{o \in \mathcal{O}_\tau} \big(p^y(o \mid e)\ (V^y(o) + C^x_{y:\tau}(o) - \mathit{Cost}(e))\big) \tag{3}$$

Let $delegate$ denote the non-deterministic function that, given a task τ, a provider y, and a contract $C^x_{y:\tau}$, returns a pair of elements: (1) an element of $\mathcal{O}_\tau \cup \{abs\}$ which depends on the effort resulting from y's decision process (obtained from Eqs. (2) and (3)), and (2) the net utility for y.

As mentioned above, we seek a global approximation for incentives. This means we must consider the effects of task delegation on all agents in the system, and to this end, we define the *social utility* as the sum of utilities obtained by delegating a task τ between all agents in the system, given a contract $C^x_{y:\tau}$:

$$sU(C^x_{y:\tau}) = \sum_{x,y \in \mathcal{A}} U^x(\widehat{o}) + \widehat{nV^y} \tag{4}$$

where $\langle \widehat{o}, \widehat{nV^y} \rangle = delegate(\tau, y, C^x_{y:\tau})$.

Finally, following [20, p. 51], we can compute the *regret* of a consumer to have chosen a contract $C^x_{y:\tau}$ from a set of contracts $\mathfrak{C}^x_{y:\tau}$ as follows.

$$\mathit{Regret}(C^x_{y:\tau}, o) = \left(\min_{\overline{C^x_{y:\tau}} \in \mathfrak{C}^x_{y:\tau}} \overline{C^x_{y:\tau}}(o) \right) - C^x_{y:\tau}(o) \tag{5}$$

The regret value is, by definition, negative. However, its value must be interpreted as an absolute value [20, p. 51].

2.4 Traditional Solutions to the Cold Start Problem

Several approaches have been proposed for solving the cold start problem in contract negotiation. For the purpose of this work we will focus on three of them.

The first approach is probably the simplest. It proposes that as starting point, one uses the minimum contract possible according to fair systems requests. In this way, by incrementing the value of the contract (in the sense of utility paid to the provider for successful task execution) every time we receive an abstention, we can converge to the minimum contract while guaranteeing providers not abstain. However, this contract is not guaranteed to maximise the consumer's expected utility.

A second approach adopts an exploration strategy, such as Boltzmann selection [6], whereby, given a set of outcomes for a task, contracts are randomly selected initially (i.e. the exploration phase), with the best observed contract being chosen after some time period has elapsed (the exploitation phase). In case the chosen contract does not guarantee participation, we can iterate the interactions and converge sooner to the minimum contract guaranteeing it. However, if the cost of the contract for the consumer is too high, this will result in high regret.

A third approach utilises sampling. This includes widely adopted techniques for solving the cold-start problem via active learning [22]. Given a society of agents $\mathcal{A}$ and the set of contracts $\mathcal{C}$, and given a simple sampling procedure [8], the problem is to determine the number of samples required to achieve some statistical accuracy requirement.

Since our goal is to identify sufficient samples to obtain non-abstaining behaviours, we can divide each unit of the search space $\mathcal{A} \times \mathcal{C}$ — assumed to be normally distributed — into one of two classes yielding either abstaining or non-abstaining output. Given the margin of error d that we consider acceptable in the estimated proportion p in the class of non-abstaining output, and given the accepted risk α that we can incur that the actual error is larger than d, then according to [8], we require n samples, computed as follows.

$$n_0 = \frac{t^2 \, p \, (1-p)}{d^2}$$

where t is the abscissa of the normal curve that cuts off an area of α at the tails. If $\frac{n_0}{|\mathcal{A} \times \mathcal{C}|}$ is negligible, n_0 is a satisfactory approximation, and thus $n = n_0$. Otherwise,

$$n = \frac{n_0}{1 + \frac{n_0}{|\mathcal{A} \times \mathcal{C}|}} \tag{6}$$

In the following we identify with $\widehat{\mathcal{C}}^d_\alpha$ the set of sample contracts given the margin of error d and the accepted risk α.

3 Global Contracts

In order to apply PAT, one must be able to compute the provider's expected utility, requiring knowledge about provider costs and success likelihoods for different effort levels. Therefore, in order to assess its own utility, the consumer should know, for each provider and for each effort level, the associated cost, as well as the probability of obtaining each outcome of the task for a given provider's effort level.

To reduce the computational effort for a consumer to explore providers' capabilities, we introduce a *representative agent* ω obtained from the providers present in the system. ω can be viewed as the simplest stereotype agent [1, 10, 4] for a given society. The idea is to improve agents' local computations by providing them with this representative agent. Although outside the scope of this work, assessing the quality of the representative agent is an important issue which can impact on other aspects of a multi-agent society, e.g., how much can an agent trust the agents in a society in which it enters for the very first time?

In what follows, without loss of generality and to simplify the presentation, we assume:

- a single task τ;
- a fixed shared cost function $Cost$ for all the providers;
- a fixed shared utility function V for all the providers;
- a fixed shared utility function U for all the consumers;
- contracts $C^x_{y:\tau}$ such that $\forall o \in \mathcal{O}_\tau, C^x_{y:\tau}(o) \in \mathbb{Z}$: moreover, we assume a strong fairness requirement for these contracts, i.e., $C^x_{y:\tau}(o_i) \gneq C^x_{y:\tau}(o_j)$ if $o_i \succ_o o_j$.

In particular, shared utility functions are widely employed in cooperative contexts [3], which is also the main focus of our research.

Therefore, the representative agent ω is one such that:

$$\forall o \in \mathcal{O}_\tau, \forall e \in \mathcal{E}_\tau, p^\omega(o \mid e) = \frac{1}{|\mathcal{A}|} \sum_{y \in \mathcal{A}} p^y(o \mid e)$$

3.1 Searching for a Suitable Contract as a Linear Problem

Considering a representative agent as the "average" provider in a given society does not entirely address the problem of identifying a suitable contract. However, by taking into account Equation (2), we can derive *bounds* for the contracts, such that values of contract below the lower bound would have the same effect as the minimum contract itself, and the same for the upper bound.

The lower bound for contracts is:

$$\forall o \in \mathcal{O}_\tau,\ C^x_{\omega:\tau}(o) \geq \left\lfloor \left(\min_{e \in \mathcal{E}_\omega} Cost^\omega(e) \right) - \left(\max_{o \in \mathcal{O}_\tau} V^\omega(o) \right) \right\rfloor \tag{7}$$

Similarly, the upper bound is:

$$\forall o \in \mathcal{O}_\tau,\ C^x_{\omega:\tau}(o) \leq \left\lceil \left(\max_{e \in \mathcal{E}_\omega} Cost^\omega(e) \right) - \left(\min_{o \in \mathcal{O}_\tau} V^\omega(o) \right) \right\rceil \tag{8}$$

Given the bounds of Equations (7) and (8), let $\mathfrak{C}^x_{\omega:\tau} \subseteq C^x_{y:\tau}$ be the set of contracts that respect them.

Recall that our aim is to identify a suitable global contract given limited knowledge of the providers, taking into account the trade-off between (i) maximising the social utility, while (ii) minimising the (absolute value of) regret for the consumer.

Concerning (i), from Equation (2) there is an inverse relationship between the likelihood of abstaining from accepting the task and the utility gained by the provider. This thus limits our search space, as we want to select a contract that is not likely to lead to an abs result. Concerning (ii), Equation (5) suggests that minimising the chosen contract is corellated with minimising the (absolute value of the) regret as well. Let us notice that this requirement does not apply to contracts in general, rather regret minimisation is enforced only in searching for a suitable contract given the representative agent. There might be situations where minimising the regret for the consumer is unnecessary: we will investigate them in future work.

Solving the following linear problem thus addresses the above two aims:

$$\min_{C^x_{\omega:\tau}} \sum_{o \in \mathcal{O}_\tau} C^x_{\omega:\tau}(o) \tag{9}$$

subject to

$$\begin{aligned} &\sum_{o \in \mathcal{O}_\tau} C^x_{\omega:\tau}(o) \sum_{e \in \mathcal{E}_\omega} p^\omega(o \mid e) \geq \\ &\quad V^\omega(abs) - \left(\sum_{o \in \mathcal{O}_\tau} \sum_{e \in \mathcal{E}_\omega} (V^\omega(o) - Cost^\omega(e)) \right) \end{aligned} \tag{10}$$

and

$$\begin{aligned} &\forall o \in \mathcal{O}_\tau \\ &\quad C^x_{\omega:\tau}(o) \geq \left(\min_{e \in \mathcal{E}_\omega} Cost^\omega(e) \right) - \left(\max_{o \in \mathcal{O}_\tau} V^\omega(o) \right), \text{and} \\ &\quad C^x_{\omega:\tau}(o) \leq \left(\max_{e \in \mathcal{E}_\omega} Cost^\omega(e) \right) - \left(\min_{o \in \mathcal{O}_\tau} V^\omega(o) \right) \end{aligned} \tag{11}$$

and

$$\forall o_i, o_j \in \mathcal{O}_\tau \text{s.t.} o_i \succ_o o_j,\ C^x_{\omega:\tau}(o_i) > C^x_{\omega:\tau}(o_j) \tag{12}$$

In particular, Equation (9) seeks to minimise regret, while Equation (10) constrains the search space to avoid abstentions. Equations (11) and (12) enforce the lower and upper bounds on the contract, as well as the fairness constraint respectively.

[7] shows how the contract which provides a solution to the linear problem is a suitable — i.e., a trade-off between social utility and regret — approximation to the best solution obtained through exhaustive search.

3.2 Sampling the Society

Deriving a representative agent is clearly a complex task. If there is no *a priori* knowledge to do so, then performing the derivation forms an instance of the cold start problem. Therefore, we can adapt the idea of simple sampling discussed in Section 2.4 to the case of continuous data [8].

Let us assume the society $\mathcal{A}$ is distributed as a normal distribution $\mathcal{N}(\mu, \sigma)$ with mean μ and standard deviation σ. Given r the acceptable relative error, and α, the risk of being mislead by the sample, the size of the sample n required is as follows.

$$n_0 = \frac{t^2\ S^2}{r^2\ \mu^2}$$

Algorithm 1 Increment Contracts

$\mathbf{contractIncrement}(C^x_{y:\tau})$

1: **Input:** $C^x_{y:\tau}$ a valid contract
2: **Output:** $\overline{C^x_{y:\tau}}$ an incremented contract
3: $\overline{C^x_{y:\tau}} := C^x_{\omega:\tau}$
4: **for** $o_i \in \langle o_1, \ldots, o_n \rangle$ s.t. $\forall j, k, j > k, o_j \succ_o o_k$ **do**
5:     **if** $o_i = o_1$ **then**
6:         $\overline{C^x_{y:\tau}}(o_i) = \overline{C^x_{y:\tau}}(o_i) + 1$
7:     **else**
8:         **if** $\overline{C^x_{y:\tau}}(o_{i-1}) = \overline{C^x_{y:\tau}}(o_i)$ **then**
9:             $\overline{C^x_{y:\tau}}(o_i) := \overline{C^x_{y:\tau}}(o_i) + 1$
10:         **end if**
11:     **end if**
12: **end for**
13: **return** $\overline{C^x_{y:\tau}}$

Here $S^2 = \frac{\sum_{i=1}^{|\mathcal{A}|}(a_i - a_\mu)}{|\mathcal{A}|-1} \simeq \sigma^2$. If $\frac{n_o}{|\mathcal{A}|}$ is appreciable, $n = \frac{n_0}{1+\frac{n_0}{|\mathcal{A}|}}$ (cf. Eq. 6); otherwise $n = n_0$ [8].

It is worth noticing that we only sample $\mathcal{A}$, while the equations in Section 2.4 sample the space $\mathcal{A} \times \mathcal{C}$. In the following, $\mathcal{A}^r_\alpha$ identifies the sampled space of agents with r relative error and α the risk of being mistaken; $\widehat{\omega}^r_\alpha$ identifies the representative agent derived from simple sampling of $\mathcal{A}^r_\alpha$.

4 Global Contract for Cold Start Problem

We now focus on addressing the cold start problem [23], viz. the problem to identify a suitable contract to be used in PAT with minimal information about providers.

4.1 Searching for Non-Abstaining Contracts

According to Equation (2), given a contract $C^x_{y:\tau}$ a provider will abstain from performing the task τ if doing so will increase its expected utility. In such cases, it is necessary to increment $C^x_{y:\tau}$ towards a "better" (to the provider) fair contract.

To this end, Algorithm 1 defines the $\mathbf{contractIncrement}$ procedure which returns the closest higher fair contract $\overline{C^x_{y:\tau}}$ of an contract $C^x_{y:\tau}$ given as input.

At line 3 of Algorithm 1, $\mathbf{contractIncrement}$ copies the value of $C^x_{y:\tau}$ to $\overline{C^x_{y:\tau}}$; at line 6 it increments the value of the contract for the worst outcome $\overline{C^x_{y:\tau}}(o_1)$. Then, for each other outcome o_i in the sequence induced by the ordering function $\succ_o$, $\mathbf{contractIncrement}$ checks if $\overline{C^x_{y:\tau}}(o_{i-1}) = \overline{C^x_{y:\tau}}(o_i)$ (line 8). If this is the case, $\overline{C^x_{y:\tau}}(o_i)$ is also incremented to ensure fairness.

The following proposition proves that there are no other contracts "smaller" than $\overline{C^x_{y:\tau}} = \mathbf{contractIncrement}(C^x_{y:\tau})$ but "greater" than $C^x_{y:\tau}$. Therefore, $\overline{C^x_{y:\tau}}$ is the closest of the contracts (in terms of increments necessary within them) that are more convenient (for the provider) than $C^x_{y:\tau}$.

Algorithm 2 Hill-Climbing Contracts

$\mathbf{hillC}(y, C^x_{y:\tau})$

1: **Input:** $y \in \mathcal{A}, C^x_{y:\tau}$ a valid contract
2: **Output:** S the number of iterations to non-abstain behaviour
3: $S := 0$
4: **while** $delegate(\tau, y, C^x_{y:\tau}) = abs$ **do**
5: $\quad S := S + 1$
6: $\quad C^x_{y:\tau} := \mathbf{contractIncrement}(C^x_{y:\tau})$
7: **end while**
8: **return** S

Proposition 1. *Given a contract* $C^x_{y:\tau}$, *and* $\overline{C^x_{y:\tau}} = \mathbf{contractIncrement}(C^x_{y:\tau})$, $\forall o \in \mathcal{O}_\tau \nexists \widehat{C^x_{y:\tau}} \in \mathfrak{C}^x_{y:\tau} \setminus \{C^x_{y:\tau}, \overline{C^x_{y:\tau}}\}$ *s.t.* $C^x_{y:\tau}(o) < \widehat{C^x_{y:\tau}(o)} < \overline{C^x_{y:\tau}}(o)$.

Proof. Let assume that $\exists \widehat{C^x_{y:\tau}}$ s.t. $C^x_{y:\tau}(o) < \widehat{C^x_{y:\tau}(o)} < \overline{C^x_{y:\tau}}(o)$ for some $o \in \mathcal{O}_\tau$.

If $o = o_1$, from line 6 of Algorithm 1 $\widehat{C^x_{y:\tau}} = \overline{C^x_{y:\tau}}$, *quod est absurdum.*

If $o = o_i, i > 1$, without loss of generality let us assume $\overline{C^x_{y:\tau}}(i) = \widehat{C^x_{y:\tau}(o_i)} + 1 = C^x_{y:\tau}(o_i) + 2$. From l. 9 of Algorithm 1, this implies that $C^x_{y:\tau}(o_i) = C^x_{y:\tau}(o_i) + 1$, *quod est absurdum.* □

While Algorithm 1 derives more convenient (for the provider) contracts, Algorithm 2 implements the sound and complete procedure $\mathbf{hillC}$ for computing the distance of a given contract $C^x_{y:\tau}$ from the closest contract which is more convenient (for the provider) than abstaining from performing the given task.

$\mathbf{hillC}$ requires as input the provider y, and a contract $C^x_{y:\tau}$. It returns the number of interactions with y needed to ensure that y will not abstain. At line 3 it initialises the variable S which stores the number of interactions with y. Such a variable is incremented (l. 5) every time the delegation process returns abs.[1] In such a case, the contract is incremented (l. 6) using the function $\mathbf{contractIncrement}$ (Algorithm 1).

The following proposition proves that $\mathbf{hillC}$ (Algorithm 2) is complete and sound.

Proposition 2. *Algorithm 2 is sound and complete.*

Proof. Immediate from Prop. 1 and Eq. (2). □

4.2 Experimental Hypotheses

The procedure $\mathbf{hillC}$ takes a contract as input: determining which contract to use first is the essence of the cold start problem. For the purpose of this work, we compare the following possible initial contracts:

[1] We admit a small abuse of notation: formally $delegate$ returns a tuple of two elements. In this case we silently assume that returns only the first element of such a tuple, namely the outcome of task τ or abs.

- $C^x_{y:\tau\downarrow\text{global}}$ s.t. $\forall o \in \mathcal{O}_\tau, C^x_{y:\tau\downarrow\text{global}}(o) = C^x_{\omega:\tau}(o)$, where $C^x_{\omega:\tau}$ is a solution to the linear problem of Eqs. (9–12). We denote $C^x_{y:\tau\downarrow\text{global}}$ as GLOBAL. In the following, to show the robustness of our approach, we assume that the capabilities of the representative agents are uniformly perturbed by up to 0.2 from the average.
- $C^x_{y:\tau\downarrow\text{globalS}}$ s.t. $\forall o \in \mathcal{O}_\tau$, $C^x_{y:\tau\downarrow\text{globalS}}(o) = C^x_{\widehat{\omega}^r_\alpha:\tau}(o)$, where $C^x_{\omega:\tau}$ is a solution to the linear problem of Eqs. (9–12). We denote $C^x_{y:\tau\downarrow\text{globalS}}$ as GLOBALsample. We considered $\alpha = 0.05$, and $d = 0.20$.

We compare these two contracts to three contracts capturing existing approaches to dealing with the cold start problem.

- $C^x_{y:\tau\downarrow\text{min}}$ s.t. $\forall o \in \mathcal{O}_\tau, C^x_{y:\tau\downarrow\text{min}}(o) = \min_{C^x_{y:\tau}} C^x_{y:\tau}(o)$ according to Eq. (7). We denote $C^x_{y:\tau\downarrow\text{min}}$ as MIN;
- $C^x_{y:\tau\downarrow\text{rand}}$ s.t. $\forall o \in \mathcal{O}_\tau, C^x_{y:\tau\downarrow\text{rand}}(o) = \mathsf{rand}$, where rand is a random number between 0 and 1 derived from an uniform distribution. We denote $C^x_{y:\tau\downarrow\text{rand}}$ as RANDOM;
- $C^x_{y:\tau\downarrow\text{S}}$ s.t. $\forall o \in \mathcal{O}_\tau, C^x_{y:\tau\downarrow\text{S}}(o) = \min_{C^x_{y:\tau} \in \widehat{\mathcal{C}}} C^x_{y:\tau}(o)$ according to Eq. (7). We denote $C^x_{y:\tau\downarrow\text{S}}$ as CONTRACTsample. We considered $\alpha = 0.05$, and $d = 0.20$.

Our experimental hypotheses are:

I1: on average, the procedure **hillC** invoked on GLOBAL and GLOBALsample (the contracts derived from the representative agent resp. without or with a sampling activity) will require a minor number of interactions to converge to a non-abstaining contract than if it is invoked on, in order, CONTRACTsample, RANDOM, **hillC**(MIN);

I2: GLOBAL and GLOBALsample are more robust with respect to changes of network structure and distribution of competencies in the network.

4.3 Experimental Settings

We ran a set of experiments to evaluate the hypotheses, as detailed below. For all experiments, we used the following base settings. $\mathcal{O}_\tau = \{o_1, o_2\}$ s.t. $o_2 \succ_o o_1$. $V(o_1) = -10$, $V(o_2) = 50$, and $V(abs) = 0$. $\mathcal{E}_\tau = \{e_1, e_2, e_3\}$ s.t. $e_3 \succ_e e_2$, $e_2 \succ_e e_1$, and $Cost(e_1) = 10$, $Cost(e_2) = 15$, $Cost(e_3) = 20$.

Our system consisted of three agent types ($\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$), described in Table 1. We considered societies of 100 agents.

Finally, we handcrafted three different distributions of competencies among agents, namely *poorly* resp. *uniformly* resp. *highly* competent societies. Table 2 details the three distributions.

For each configuration (distribution of competencies), we generated 10 different societies (s0–s9) and evaluated each of them 50 times.

4.4 Experimental Evaluation

Figures 1, 2, and 3 qualitatively depict the results of our experiments: each figure refers to a single distribution of competencies. Figure 1 illustrates our results for poor competence societies, 2 for uniform competence, and 3 for high competence.

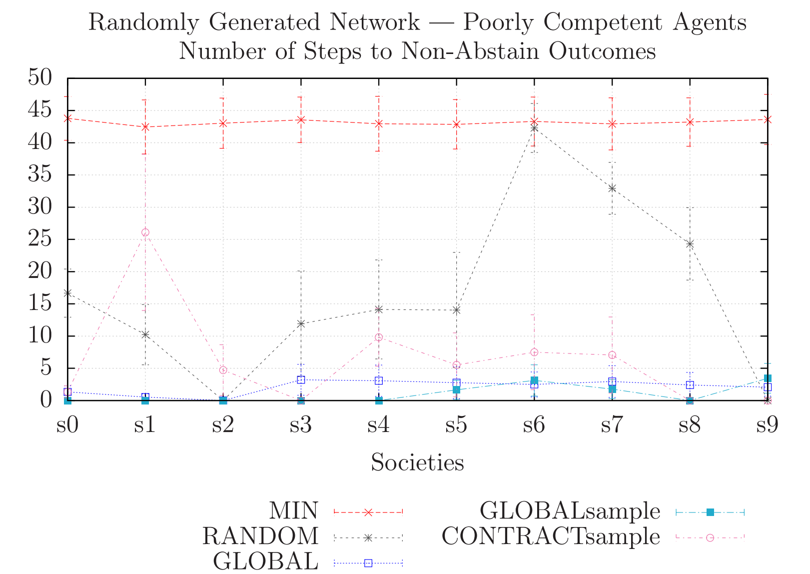

Fig. 1: Average steps and standard deviation necessary to find a contract which avoids abstaining: poorly competent agents

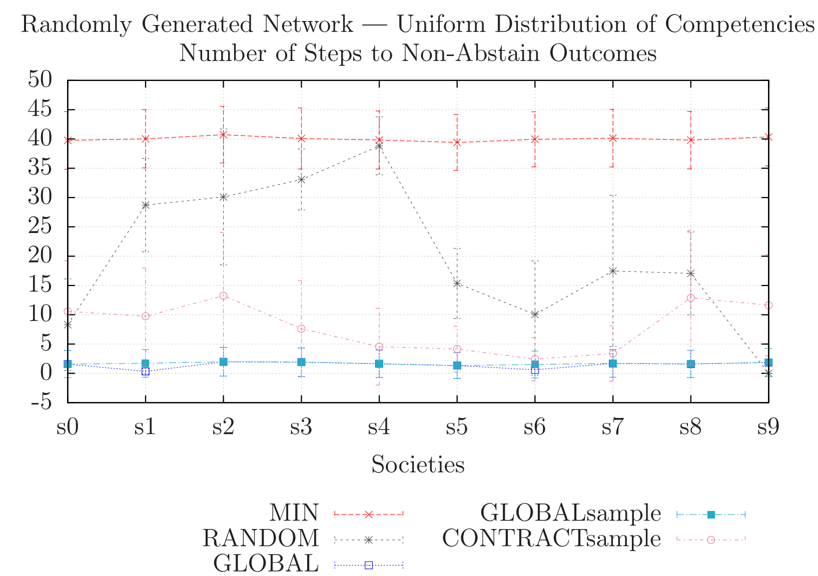

Fig. 2: Average steps and standard deviation necessary to find a contract which avoids abstaining: uniform distribution of competencies

	$\mathcal{G}_1$	$\mathcal{G}_2$	$\mathcal{G}_3$
$p^g(o_1 \mid e_1)$	0.80	0.75	0.70
$p^g(o_1 \mid e_2)$	0.60	0.55	0.50
$p^g(o_1 \mid e_3)$	0.40	0.40	0.20
$p^g(o_2 \mid e_1)$	0.20	0.25	0.30
$p^g(o_2 \mid e_2)$	0.40	0.45	0.50
$p^g(o_2 \mid e_3)$	0.40	0.60	0.80

Table 1: Agents grouped by competencies

Competence	$p(a \in \mathcal{G}_1)$	$p(a \in \mathcal{G}_2)$	$p(a \in \mathcal{G}_3)$
Poor	0.6	0.3	0.1
Uniform	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
High	0.1	0.3	0.6

Table 2: Distribution of competencies in the network

For each configuration, for each society, Figures 1, 2, and 3 show average and standard deviation — over the 50 explorations for each society, and over the agents in the society — of the number of steps needed to find a contract with no abstaining results when the starting contract is GLOBAL, GLOBALsample, MIN, RANDOM, or CONTRACTsample. Although it is not always the case that these values are normally distributed, we chose to represent average and standard deviation for qualitative purposes.

These results have been proven statistically significant using the Wilcoxon Signed-Rank Test (WSRT) [27] ($p < 0.01$). From Figures 1, 2, and 3 it is clear that our hypotheses are satisfied. With respect to hypothesis *I1*, the average of **hillC**(MIN) is the highest in any configuration and for any society. On average, values of **hillC**(RANDOM) are always smaller than **hillC**(MIN), and **hillC**(CONTRACTsample) is smaller that both of them. We also perturbed the representative agent by up to 0.2 to evaluate the resilience of the generated contract. In such a situation, **hillC**(GLOBAL) is almost always the smallest, and barely distinguishable from **hillC**(GLOBALsample). This is particularly true when the competencies are distributed uniformly, Figure 2, or society is highly competent (Figure 3).

Regarding hypothesis *I2*, it is worth noting that the distribution of competencies has an effect on **hillC**(MIN) which varies in the range $[32, 47]$, where 32 is the minimum in the case of highly competent societies, and 47 is the maximum in poorly competent societies: the poorer the agents in terms of competencies, the higher the (average) value of **hillC**(MIN). In contrast, both **hillC**(GLOBAL) and **hillC**(GLOBALsample) re-

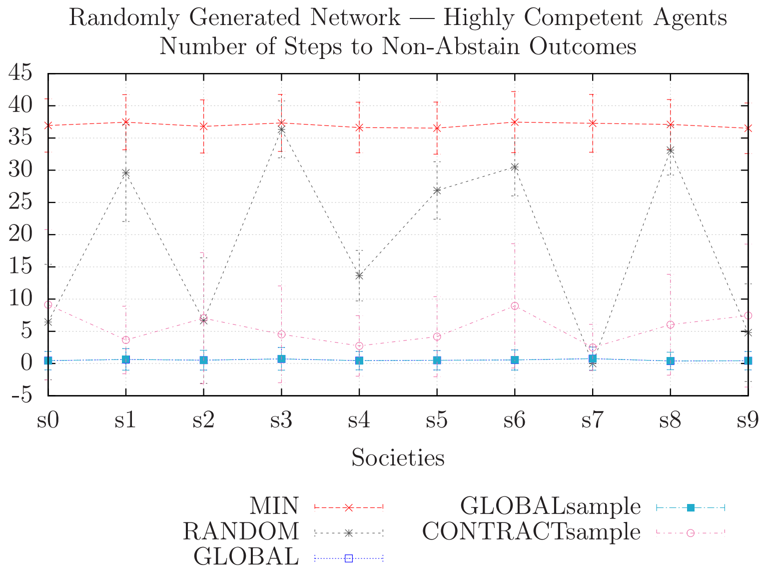

Fig. 3: Average steps and standard deviation necessary to find a contract which avoids abstaining: highly competent agents

turns values that are always in the range $[0, 8]$, independent of the configuration and the societies. This suggests that *I2* is also verified by this set of experiments.

In order to test the robustness of our approach, we also perturbed the competencies of the representative agent by adding uniform noise in intervals of 0.05 between 0 and 0.25 for each $p^g(o \mid e)$. In 53% of cases — 80% between uniformly distributed and highly competent agents — they lead to non-significant results ($p > 0.01$) according to the Kruskal-Wallis test [15]. Only in the case of poorly competent agents were these perturbations always significant. This supports our previous analysis regarding the robustness of our approach and the good performance obtained.

5 Conclusions and Future Work

In this paper, we propose techniques to identifying an approximate contract, to be used by consumers before they have obtained sufficient information to craft a specific contract for interactions with providers within the system. This contract provides a trade-off between social utility and regret, and is identified by solving a linear optimisation problem. Our work addresses an instance of the cold-start problem. We evaluated our approach empirically, comparing it to existing cold-start mitigation techniques, and found that our heuristic GLOBAL is robust over different network topologies and provider competencies; furthermore, in its GLOBALsample version, it is also computationally efficient, sampling only a small subset of agents within the system. Finally, it allows agents to converge to contracts which minimise the level of abstention with fewer interactions than heuristics derived from existing cold-start mitigation techniques.

In the empirical evaluation presented in this paper we considered binary task outcomes and discrete domains. Though this covers many situations where an agent is concerned only with the success or failure of the task, there are situations where a more fine-grained set of outcomes should be considered. For instance, in the context of information sharing, an *information consumer* might pay more for higher quality data from *information providers*. To represent these situations we plan to adapt the formalisms presented in [2, 25], where the authors discuss a framework for information sharing with different levels of quality within a multi-agent system. Moreover, our approach can be integrated with Carmel and Markovitch [6] to improve the effectiveness of their strategies for explorations.

Moreover, considering (potentially infinite) sets of outcomes and efforts would allow us to integrate our work with trust and reputation systems (e.g., [12, 26, 13]). We envision that an agent, on entering the system, would use the mechanisms described in the current work, but as it gains experience, would transition to utilising its trust and reputation mechanism to identify optimal contracts (c.f., [5]).

Acknowledgments

We thank the anonymous reviewers for their helpful comments.

Research was sponsored by US Army Research laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

This work was performed using the Maxwell High Performance Computing Cluster of the University of Aberdeen IT Service (`http://www.abdn.ac.uk/staffnet/research/research-computing`), provided by Dell Inc. and supported by Alces Software.

References

1. Beckett, N.E., Park, B.: Use of Category Versus Individuating Information: Making Base Rates Salient. Personality and Social Psychology Bulletin 21(1), 21–31 (Jan 1995), `http://psp.sagepub.com/content/21/1/21.refs`
2. Bisdikian, C., Tang, Y., Cerutti, F., Oren, N.: A Framework for Using Trust to Assess Risk in Information Sharing. In: Proocedings of 2nd International Conference on Agreement Technologies – AT 2013. pp. 135–149. Springer Berlin Heidelberg (2013), `http://link.springer.com/content/pdf/10.1007%2F978-3-642-39860-5_11.pdf`
3. Boella, G., Damiano, R., Lesmo, L.: Cooperation and group utility. In: Jennings, N., Lesprance, Y. (eds.) Intelligent Agents VI. Agent Theories, Architectures, and Languages,

Lecture Notes in Computer Science, vol. 1757, pp. 319–333. Springer Berlin Heidelberg (2000)
4. Burnett, C., Norman, T.J., Sycara, K.: Bootstrapping trust evaluations through stereotypes. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems. pp. 241–248 (May 2010), http://dl.acm.org/citation.cfm?id=1838206.1838240
5. Burnett, C., Norman, T.J., Sycara, K.: Trust decision-making in multi-agent systems. In: Proceedings of the Twenty-Second international joint conference on Artificial Intelligence — IJCAI'11. pp. 115–120. AAAI Press (Jul 2011), http://dl.acm.org/citation.cfm?id=2283396.2283417
6. Carmel, D., Markovitch, S.: Exploration strategies for model-based learning in multi-agent systems: Exploration strategies. Autonomous Agents and Multi-agent systems 2(2), 141–172 (1999)
7. Cerutti, F., Oren, N., Burnett, C.: Global Approximations for Principal Agent Theory. In: Proceedings of the 14th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2015). pp. 1845—-1846 (2015), http://aamas2015.com/en/AAMAS_2015_USB/aamas/p1845.pdf
8. Cochran, W.G.: Sampling Techniques. John Wiley & Sons, Inc., third edn. (1977)
9. Grossman, S.J., Hart, O.D.: An analysis of the principal-agent problem. Econometrica: Journal of the Econometric Society 51(1), 7–45 (1983)
10. Hilton, J.L., von Hippel, W.: Stereotypes. Annual review of psychology 47, 237–271 (Jan 1996), http://www.annualreviews.org/doi/abs/10.1146/annurev.psych.47.1.237
11. Holmstrom, B.R., Tirole, J.: The theory of the firm. Handbook of industrial organization 1, 61–133 (1989)
12. Jøsang, A., Ismail, R.: The Beta Reputation System. In: Proc. 15th Bled Electronic Commerce Conf. (2002)
13. Kamvar, S.D., Schlosser, M.T., Garcia-Molina, H.: The eigentrust algorithm for reputation management in p2p networks. In: Proceedings of the 12th international conference on World Wide Web. pp. 640–651. ACM (2003)
14. Karimi, R., Freudenthaler, C., Nanopoulos, A., Schmidt-Thieme, L.: Active learning for aspect model in recommender systems. In: 2011 IEEE Symposium on Computational Intelligence and Data Mining. pp. 162–167 (2011), http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5949431
15. Kruskal, W.H., Wallis, W.A.: Use of Ranks in One-Criterion Variance Analysis. Journal of the American Statistical Association 47(260), pp. 583–621 (1952), http://www.jstor.org/stable/2280779
16. Marsa-Maestre, I., Klein, M., Jonker, C.M., Aydoan, R.: From problems to protocols: Towards a negotiation handbook. Decision Support Systems 60, 39–54 (Apr 2014), http://www.sciencedirect.com/science/article/pii/S016792361300167X
17. Meneguzzi, F., Miles, S., Luck, M., Holt, C., Smith, M.: Electronic contracting in aircraft aftercare: a case study. In: Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems — AAMAS 08. pp. 63–70. International Foundation for Autonomous Agents and Multiagent Systems (May 2008), http://dl.acm.org/citation.cfm?id=1402795.1402807
18. Miles, S., Groth, P., Oren, N., Luck, M.: Handling Mitigating Circumstances for Electronic Contracts. In: Proceedings of the 7th European Workshop on Multi-Agent Systems (2009)
19. Miller, G.J., Whitford, A.B.: Trust and Incentives in Principal-Agent Negotiations: The 'Insurance/Incentive Trade-Off'. Journal of Theoretical Politics 14(2), 231–267 (Apr 2002), http://jtp.sagepub.com/content/14/2/231.abstract

20. Peterson, M.: An Introduction to Decision Theory. Cambridge University Press (2009), http://dx.doi.org/10.1017/CBO9780511800917
21. Popescul, A., Pennock, D.M., Lawrence, S.: Probabilistic models for unified collaborative and content-based recommendation in sparse-data environments. In: Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence. pp. 437–444. Morgan Kaufmann Publishers Inc. (2001)
22. Provost, F., Melville, P., Saar-Tsechansky, M.: Data acquisition and cost-effective predictive modeling. In: Proceedings of the ninth international conference on Electronic commerce. p. 389. New York, New York, USA (2007), http://dl.acm.org/citation.cfm?id=1282100.1282172
23. Schein, A.I., Popescul, A., Ungar, L.H., Pennock, D.M.: Methods and metrics for cold-start recommendations. In: Proceedings of the 25th international ACM SIGIR conference on Research and development in information retrieval. p. 253. New York, New York, USA (Aug 2002), http://dl.acm.org/citation.cfm?id=564376.564421
24. Schmidt, S., Steele, R., Dillon, T.S., Chang, E.: Fuzzy trust evaluation and credibility development in multi-agent systems. Applied Soft Computing 7(2), 492 – 505 (2007), http://www.sciencedirect.com/science/article/pii/S1568494606000755
25. Tang, Y., Cerutti, F., Oren, N., Bisdikian, C.: Reasoning about the Impacts of Information Sharing. Information Systems Frontiers Journal pp. 1–18 (2014), http://link.springer.com/article/10.1007%2Fs10796-014-9521-6
26. Teacy, W.T.L., Patel, J., Jennings, N.R., Luck, M.: TRAVOS: Trust and Reputation in the Context of Inaccurate Information Sources. Autonomous Agents and Multi-Agent Systems 12(2), 183–198 (Feb 2006), http://link.springer.com/10.1007/s10458-006-5952-x
27. Wilcoxon, F.: Individual Comparisons by Ranking Methods. Biometrics Bulletin 1(6), 80–83 (1945)