

Global Approximations for Principal Agent Theory

(Extended Abstract)

Federico Cerutti
University of Aberdeen, UK
f.cerutti@abdn.ac.uk

Nir Oren
University of Aberdeen, UK
n.oren@abdn.ac.uk

Christopher Burnett
University of Aberdeen, UK
cburnett@abdn.ac.uk

ABSTRACT

Principal Agent Theory (PAT) seeks to identify incentives and sanctions that a consumer should offer a producer as part of a contract in order to maximise the former’s utility. However, identifying optimal contracts in large systems is difficult, particularly when little information is available about producer competencies. In this work we propose that a global contract 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. Finally, we suggest how our work can be integrated with existing work on multi-agent systems.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

Keywords

Principal Agent Theory; Cold Start Problem

1. INTRODUCTION

In order to obtain desirable behaviour when interacting with others, an agent may be required to provide incentives. Principal Agent Theory (PAT) [2] aims to determine the optimal level of incentives — in the form of rewards and penalties — that an agent (the *principal*) must provide to others (the *providers*) in order to have these act so as to maximise the principal’s utility. PAT requires several inputs which are hard for agents to obtain, and is computationally expensive. In this paper, we seek to overcome these restrictions through approximation.

When specifying contractual incentives to those they have not dealt with before, humans often resort to general cultural or legal foundations. Building on this intuition, suggests that without additional information, a heuristic can be used to specify incentives *for all interactions within the system*. This heuristic builds on information which we assume is *a priori* available to the agent, or to the designer of the MAS.

This work makes two contributions. First, we describe a heuristic for determining optimal global incentive values. Second, we provide an analysis of the utility/incentive space, highlighting some of its properties.

Informally, our heuristic computes a contract based on an average individual provider derived from all agents in the system. In the

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

remainder of this paper, we describe how contracts are computed, and evaluate our approach, following which we conclude.

2. GLOBAL CONTRACT

Following [1], we assume a *society* of agents $\mathcal{A} = \{x, y, \dots\}$ 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* $y \in \mathcal{A}$ undertake the task on its behalf. Given $\tau \in T$, let $\mathcal{O}_\tau = \{o_1, o_2\}$ denote the set of possible *outcomes* for task τ , such that $o_2 \succ_o o_1$ — i.e., o_2 is *better* than o_1 . We assume that all agent share the same task evaluation criteria as well as the same ordering function.

In asking a provider to execute a task, the consumer *delegates* it to the provider. This delegation 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). For a task τ , the outcome can be either an element of \mathcal{O}_τ , or *abs*. The latter denotes an abstention from executing the task by the provider. Given this, the utility gained by the consumer is computed by the function $U : \mathcal{O}_\tau \cup \{abs\} \mapsto \mathbb{R}$, while the provider gains utility $V : \mathcal{O}_\tau \cup \{abs\} \mapsto \mathbb{R}$.

The task provider has autonomy in selecting the method by which a task will be carried out. In particular, let $\mathcal{E}_\tau = \{e_1, e_2, e_3\}$ denote the set of *effort levels* for task τ , such that $e_3 \succ_e e_2 \succ_e e_1$ — i.e., the highest effort is e_3 . Each effort has an associated cost determined by the function $Cost : \mathcal{E}_\tau \mapsto \mathbb{R}$. The effort also has an impact on the probability distribution over \mathcal{O}_τ : $\forall o \in \mathcal{O}_\tau, \forall e \in \mathcal{E}_\tau, p^y(o | e)$ represents the probability that the producer y will achieve the outcome o using the effort e .

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

In what follows we assume a *fair system*, viz. (1) the better the outcome of a task, the greater the utility that both the provider and the consumer receive; (2) the utility gained from abstaining cannot be strictly smaller than the utility gained from *any* outcome; (3) the higher the effort, the higher the associated cost to the provider; (4) given the highest outcome, the higher the effort, the higher the probability to achieve such an outcome; and (5) the better the outcome, the higher the compensation to the provider according to the contract.

We assume that each provider rationally decides whether or not to accept a contract depending on its expected utility [3].

Finally, following [3, p. 51], we can compute the *regret* of a consumer in having chosen a contract $C_{y:\tau}^x$ from a set of contracts $\mathcal{C}_{y:\tau}^x$.

The representative agent ω for this society is one such that:

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

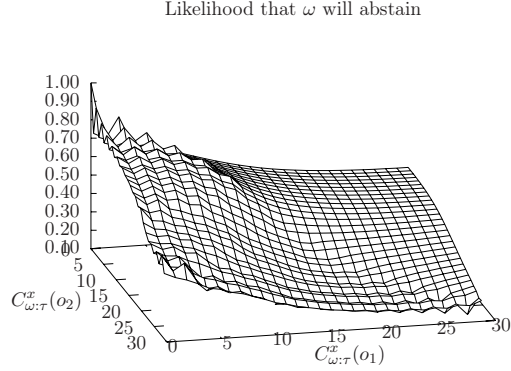


Figure 1: Brute force search for a global contract: likelihood of abstaining

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

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

$$\min_{C_{\omega:\tau}^x} \sum_{o \in \mathcal{O}_\tau} C_{\omega:\tau}^x(o) \quad (1)$$

subject to

$$\begin{aligned} & \sum_{o \in \mathcal{O}_\tau} C_{\omega:\tau}^x(o) \sum_{e \in \mathcal{E}_\omega} p^\omega(o | e) \geq \\ & V^\omega(abs) - \left(\sum_{o \in \mathcal{O}_\tau} \sum_{e \in \mathcal{E}_\omega} (V^\omega(o) - Cost^\omega(e)) \right) \\ & \forall o \in \mathcal{O}_\tau \\ & C_{\omega:\tau}^x(o) \geq (\min_{e \in \mathcal{E}_\omega} Cost^\omega(e)) - (\max_{o \in \mathcal{O}_\tau} V^\omega(o)), \\ & C_{\omega:\tau}^x(o) \leq (\max_{e \in \mathcal{E}_\omega} Cost^\omega(e)) - (\min_{o \in \mathcal{O}_\tau} V^\omega(o)) \\ & \forall o_i, o_j \in \mathcal{O}_\tau \text{ s.t. } o_i \succ_o o_j, C_{\omega:\tau}^x(o_i) > C_{\omega:\tau}^x(o_j) \end{aligned}$$

3. EVALUATION

To evaluate the effectiveness of using the global contract, we explored the social utility space for a system of 15 agents which are assembled into three homogeneous groups ($\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$) such that $\forall i \in 1 \dots 3, p^{g \in \mathcal{G}_i}(o_1 | e_1) = \langle 0.8, 0.75, 0.7 \rangle; p^{g \in \mathcal{G}_i}(o_1 | e_2) = \langle 0.6, 0.55, 0.5 \rangle; p^{g \in \mathcal{G}_i}(o_1 | e_3) = \langle 0.6, 0.4, 0.2 \rangle; p^{g \in \mathcal{G}_i}(o_2 | e_1) = \langle 0.2, 0.25, 0.3 \rangle; p^{g \in \mathcal{G}_i}(o_2 | e_2) = \langle 0.4, 0.45, 0.5 \rangle; p^{g \in \mathcal{G}_i}(o_2 | e_3) = \langle 0.4, 0.6, 0.8 \rangle. V(o_1) = -10, V(o_2) = 10, and $V(abs) = -9. Cost(e_1) = 10, Cost(e_2) = 15, Cost(e_3) = 20. The likelihood of success for each effort level for each agent were randomly perturbed from those of its base type by up to 0.2 in each direction, to reflect the differing competences of different providers. We then evaluated the social utility for the system averaged over 100 runs, where during each run, all agents in the system acted as principals and providers.$$

Figures 1 and 2 qualitatively summarise the likelihood of provider abstention and regret for different contract values. In all of these figures, the x and y axes specify the utility transferred from the consumer to the provider for each outcome. There is a clear inverse relationship between the likelihood of abstaining and the value of the contract. Figure 2, depicts the regret of a consumer. Clearly, highly positive contracts (i.e., ones in which a provider is paid regardless of contract outcomes), result in high regret (in absolute value) for the consumer. A good contract is one which trades off

Regret when the likelihood that ω will abstain is less than 0.5




Figure 2: Brute-force search for a global contract: regret in generally non-abstaining cases

these parameters, and from these figures, appears to be located in the neighbourhood of $C_{\omega:\tau}^x(o_1) = 13, C_{\omega:\tau}^x(o_2) = 13$. Therefore, with our experimental settings, the agent will abstain until a contract with a payment of at least 13 units is reached.

Once a representative agent has been identified, and the optimisation problem described in Eq. 1 solved, we obtain a contract specifying the following rewards for different outcomes: $C_{\omega:\tau}^x(o_1) = 13$; and $C_{\omega:\tau}^x(o_2) = 14$. This satisfies the first requirement for a fair system; we considered integer contracts only.

4. CONCLUSIONS

In this paper we proposed a global approximation to PAT, specifying how rewards and penalties within a contract can be identified which should govern all agent interactions within a system in order to maximise global utility. We also provided a brief analysis of the utility space for a representative scenario. There are several interesting avenues of future work, including integrating our results with a trust and reputation system.

Acknowledgments

This 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.

REFERENCES

- [1] C. Burnett, T. J. Norman, and K. Sycara. Trust decision-making in multi-agent systems. In *Proceedings of IJCAI'11*, pages 115–120, 2011.
- [2] S. J. Grossman and O. D. Hart. An analysis of the principal-agent problem. *Econometrica: Journal of the Econometric Society*, 51(1):7–45, 1983.
- [3] M. Peterson. *An Introduction to Decision Theory*. Cambridge University Press, 2009.