Designing Efficient Online Trading Systems

Ryan Porter*

Yoav Shoham*

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Economics, Design

Keywords

Mechanism Design, Multiagent Systems

1. INTRODUCTION

In many online domains, agents share goods or services that are both cheap to provide and valuable to receive. Examples include ratings in a recommender system, forwarding a message to node closer to its destination in a mobile ad-hoc network, and uploading a file in a P2P system. Existing systems in these domains often rely on agents to voluntarily provide goods. Despite the low cost of doing so, many agents instead choose to free-ride, leading to a loss in social efficiency. We present an initial step towards addressing this problem in the context of an abstract, extremely simplified model, where we identify more efficient mechanisms, including one that is provably optimal. Of course, we are not the first to consider this problem, or to offer a solution (see, e.g., [1] for recommender systems, [3] for ad-hoc networks, and [2] and [4] for P2P systems).

2. FORMULATION

We consider one-shot interactions between two agents: a_1 and a_2. Each agent can provide a good that the other desires with independent probability p. Each agent a_i is characterized by a type $\theta_i = \langle v_i, c_i \rangle \in \Theta$. An agent gains a value of v_i if it receives a desired good from the other agent, and incurs a cost of c_i if it provides a good. The latter variable captures both the cost of delivering the good and the value for contributing to the system, and thus can be either positive or negative.

For expositional purposes, we will only consider two possible types: $\Theta = (\theta_i^a, \theta_i^s)$. For both types, $v_i = v$ and $|c_i| = c$, where v and c are positive constants. The difference is that $c_i = c$ for θ_i^s (the *selfish* type) and $c_i = -c$ for θ_i^a (the *altruistic* type). Independently for each agent, $\theta_i = \theta_i^s$ with probability P_s.

The set of possible outcomes contains nine different trading possibilities: $O = \{(\leftarrow_\perp, \rightarrow_\perp), (\leftarrow_\perp, \rightarrow_1), (\leftarrow_\perp, \rightarrow_{1,2}), (\leftarrow_2, \rightarrow_\perp), (\leftarrow_2, \rightarrow_1), (\leftarrow_2, \rightarrow_{1,2}), (\leftarrow_{1,2}, \rightarrow_\perp), (\leftarrow_{1,2}, \rightarrow_1), (\leftarrow_{1,2}, \rightarrow_{1,2})\}$. The left and right arrows represent an good provided by agent 2 to 1 and by agent 1 to 2, respectively. The subscript on the arrow denotes the agents that must be able to provide a good that the other agent desires in order for the transfer to occur, with "$\perp$" denoting that the transfer never occurs. The actual transaction that occurs comes from the set $T = \{-, \leftarrow, \rightarrow, \leftrightarrow\}$, to represent, respectively: no transfers, a transfer from agent 2 to 1 only, a transfer from agent 1 to 2 only, and a transfer in both directions. An agent's valuation for a transfer is denoted by $\overline{u}_i(t, \theta_i)$.

An outcome $o \in O$, combined with p, induces a lottery over transactions. The utility function for an agent, $u_i : O \times \Theta \rightarrow \Re$, maps each outcome and agent type to its expected valuation for the transaction.

The mechanism defines the protocol for interaction between the agents and the center that culminates with the selection of an outcome. It is formally defined by a tuple $\Gamma = (A, g(\cdot))$, where A is the action space of each agent, and $g : A^2 \rightarrow O$ maps the actions of both agents to an outcome $o \in O$. A key aspect of our setting is that we are only investigating what can be achieved without the use of monetary payments.

A mechanism Γ, combined with Θ, P_s, and $u_i(\cdot)$, induces a Bayesian game, in which each agent selects a (mixed) strategy $s_i : \Theta \rightarrow \Delta A$, that maps each possible type to a distribution over actions. Let $s_i(\theta_i, a)$ denote the probability assigned to a by $s_i(\theta_i)$, and $s_i(\theta_i) = a$ denote a pure strategy. Given a strategy profile, the expected utility of a_i is: $EU_i(s) = \sum_{\theta_i, \theta_{-i} \in \Theta} \sum_{a_i, a_{-i} \in A} Pr(\theta_i) \cdot Pr(\theta_{-i}) \cdot s_i(\theta_i, a_i) \cdot s_{-i}(\theta_{-i}, a_{-i}) \cdot u_i(g(a_i, a_{-i}), \theta_i)$.

The equilibrium concept we use is that of *Bayes-Nash equilibrium* (BNE). A strategy profile $s^* = (s_1^*, s_2^*)$ is a BNE of Γ if $\forall i, s_i', EU_i(s_i^*, s_{-i}^*) \geq EU_i(s_i', s_{-i}^*)$. Due to the symmetry of the setting, we consider only symmetric equilibria, and represent them using a single s_i^*.

*Computer Science Department, Stanford University, {rwporter,shoham}@stanford.edu. This work was supported in part by DARPA grant F30602-00-2-0598.

EC'04, May 17–20, 2004, New York, New York, USA.
ACM 1-58113-711-0/04/0005.

A *social choice function* (SCF) $f : \Theta^2 \rightarrow O$ maps every profile of agent types to an outcome, and Γ *implements* $f(\cdot)$ *in BNE* if there exists a BNE s^* of Γ such that $\forall \theta, g(s^*(\theta)) = f(\theta)$. We aim to implement a SCF that maximizes the expected number of transfers.

Two assumptions about our setting impose constraints on the mechanism. First, we assume that an agent can reject the resulting transaction, and thus require that the SCF $f(\cdot)$ we implement satisfy *ex post individual rationality* (ex post IR), which holds if $\forall \theta, i, \nexists t \in T, (Prob(t|f(\theta)) > 0) \wedge (\overline{u}_i(t, \theta_i) < 0)$. Second, we assume that an agent can undetectably "sabotage" its chances of being able to provide a good that the other agent desires. We model this restriction by constraining the mechanism to include a "free-rider variant" of each action, which allows an agent to play the action without the possibility of providing a good, and without preventing the receipt of a good from an agent who is willing to unconditionally provide one.

3. RESULTS

We begin by examining mechanism Γ_{SF}, which has two actions, share (S) and free-ride (F), and induces the game specified in Table 1. In our setting, this is the most basic game one could examine: one action allows sharing, while the other is the "free-rider variant". Since it is a strictly dominant strategy for altruistic agents to play S and for selfish agents to play F, the strategy $s_i^*(\theta_i^a) = S$ and $s_i^*(\theta_i^s) = F$ forms the unique BNE.

	S	F
S	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$
F	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$

Table 1: Utilities for the row player a_i, plus outcomes, in the game induced by mechanism Γ_{SF}.

3.1 Step I: Promoting Trades

In order to induce selfish agents not to free-ride, we add the action Trade (T) to create mechanism Γ_{STF} (see Table 2). This action is similar to F; however, when both agents play T, and when both agents can provide a good to the other, a trade occurs.

	S	T	F
S	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$
T	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	$(v_i - c_i) \cdot p^2$ $(\leftarrow_{1,2}, \rightarrow_{1,2})$	0 $(\leftarrow_\perp, \rightarrow_\perp)$
F	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$

Table 2: Game induced by mechanism Γ_{STF}.

It is now a weakly dominant strategy for a selfish agent to play T. However, S is no longer dominant for an altruistic agent, since playing T instead yields a chance of receiving a good from an opponent who plays T. Additionally, even when agents play according to the strategy $s_i^*(\theta_i^a) = S$ and $s_i^*(\theta_i^s) = T$ (which is a BNE when $p \leq \frac{c}{(v+c) \cdot P_s}$), the expected number of transfers is not maximized.

3.2 Step II: Maximizing Transfers

We address these problems in Γ_{SCTF} (see Table 3) by adding a fourth action, "Conditional-Share" (C), which allows an agent to provide a good whenever possible, but demand a good in return if he provides one to an agent capable of returning the favor. Of course, if the other agent free-rides, then an agent who plays C will provide a good if possible while never receiving one, because the other agent appears to simply have no good to provide. However, an opponent who plays T must, if possible, provide a good in return for receiving one.

	S	C	T	F
S	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$
C	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$v_i \cdot p^2 - c_i \cdot p$ $(\leftarrow_{1,2}, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$
T	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	$v_i \cdot p - c_i \cdot p^2$ $(\leftarrow_2, \rightarrow_{1,2})$	$(v_i - c_i) \cdot p^2$ $(\leftarrow_{1,2}, \rightarrow_{1,2})$	0 $(\leftarrow_\perp, \rightarrow_\perp)$
F	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$

Table 3: Game induced by mechanism Γ_{SCTF}.

Now, it is a weakly dominant strategy for an altruistic agent to play C. Unfortunately, we may have undone the work of adding T, since T is no longer dominant for selfish agents, who instead must trade off the possibility of trading with a selfish agent against the possibility of unnecessarily providing a good to an altruistic agent. However, if $P_s \geq \frac{c}{v}$ (which we consider likely to hold in our motivating examples), the strategy $s_i^*(\theta_i^a) = C$ and $s_i^*(\theta_i^s) = T$ is a BNE. Thus, Γ_{SCTF} implements the SCF $f_{SCTF}(\theta_i^a, \theta_i^a) = (\leftarrow_2, \rightarrow_1)$, $f_{SCTF}(\theta_i^a, \theta_i^s) = (\leftarrow_{1,2}, \rightarrow_1)$, $f_{SCTF}(\theta_i^s, \theta_i^a) = (\leftarrow_2, \rightarrow_{1,2})$, $f_{SCTF}(\theta_i^s, \theta_i^s) = (\leftarrow_{1,2}, \rightarrow_{1,2})$. We can then show that this is optimal, even absent the "free-rider variant" constraint.

THEOREM 1. *There does not exist a mechanism that implements a SCF that satisfies ex post IR and that yields a greater expected number of transfers than* f_{SCTF}.

In the full paper, we prove this result. We also characterizes all BNE for the games induced by Γ_{STF} and Γ_{SCTF}, as a function of p, P_s, c, and v, and show why the existence of other, less desirable equilibria in the last game is unavoidable in optimal mechanisms.

4. REFERENCES

[1] Chris Avery, Paul Resnick, and Richard Zeckhauser, *The market for evaluations*, American Economic Review **89** (1999), 564–584.

[2] C. Buragohain, D. Agrawal, and S. Suri, *A game theoretic framework for incentives in p2p systems*, P2P2003, 2003.

[3] J. Crowcroft, R. Gibbens, F. Kelly, and S. Ostring, *Modelling incentives for collaboration in mobile ad hoc networks*, WiOpt'03, 2003.

[4] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge, *Incentives for sharing in peer-to-peer networks*, WELCOM'01, 2001.