Efficiency and Fairness in Team Search with Self-Interested Agents

Igor Rochlin, Yonatan Aumann, David Sarne and Luba Golosman
Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel

ABSTRACT

We consider team-work settings where individual agents incur costs on behalf of the team. In such settings it is frequently the custom to reimburse agents for the costs they incur (at least in part) in order to promote fairness. We show, however, that when agents are self-interested such reimbursement can result in degradation in efficiency - at times severe degradation. We thus study the relationship between efficiency and fairness in such settings, distinguishing between ex-ante and ex-post fairness. First, we analyze reimbursement policies that reimburse solely based on purchase receipts (as is customary), and show that with such policies the degradation in both efficiency and fairness can be unbounded. We thus introduce two other families of reimbursement policies. The first family guarantees optimal efficiency and ex-ante fairness, but not ex-post fairness. The second family improves (at times) on ex-post fairness, but at the expense of efficiency, thus providing a tradeoff between the two.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Economics, Algorithms

Keywords

Multi-Agent Exploration, Cooperation, Teamwork, Economically-Motivated Agents

1. INTRODUCTION

Cooperation and team work are key to many multi-agent systems. Team work has many advantages, but perhaps the most fundamental is *efficiency* - operating as a team, agents can benefit from the work performed by others, eliminating duplicate work and costs. Thus, a team of thirsty people in a desert can collectively dig a single well, rather than each digging its own; a group of medical students can all use the expensive book purchased by one, rather than each having to purchase it separately; and a team of shopbots can split the search space amongst themselves - and then share the information, rather than each having to search the entire space. Thus, working as a team can significantly reduce costs, both collective and individual. At the same time, while agents - even self-interested agents - are frequently accepting that others benefit from their work, they also expect that the distribution of costs amongst the team members be "fair", in the sense that costs incurred are (more or less) evenly distributed among the members. Thus, continuing the above example, the student who purchased the expensive book expects to be reimbursed for her costs. However, these two objectives - efficiency and fairness, can be at odds - as demonstrated by the following example.

Consider again the expensive book example. Suppose there are twenty medical students all of which require some book, say "The Adult and Pediatric Spine" (ISBN 0781735491). One copy of the book is sufficient for all students, so only one of them, say Alice, buys the book. For fairness, the students agree that they will reimburse Alice for their share ($\frac{19}{20}$) of the cost, whatever that will turn out to be. "The Adult and Pediatric Spine" lists for \$360, but can also be found used, with some effort, for less than \$70 (these are the true figures). However, paying only 5% of the cost (i.e. at most \$18), Alice has little incentive to invest the time and effort to get a better deal, and purchases the book at full price at the university bookshop. The team in total thus paid \$360. If, on the other hand, fairness would not have been required and Alice would incur the entire the cost of the book, she would invest some effort, say worth \$20, in getting the \$70 deal and the total cost would be \$90. Thus, the requirement for fairness has reduced the overall efficiency by a factor of four. In this case, fairness damaged the team as a whole, but at least benefited Alice. But, there are also cases where insisting on fairness can damage all. Suppose that there are twenty expensive books to purchase, and each student is tasked with getting one book. The students also agree to split all costs evenly - based on purchase receipts - as is customary. As in the previous case, the result is that each student has little incentive to invest any effort in getting a good deal on its assigned book, and the team ends up paying full price on *all* books. If, on the other hand, each student would fully pay for its assigned book, then they may end up paying different amounts - which could be considered as unfair - but each student would invest the time in getting a good deal, and the resulting cost would be lower for each and every student.

In this paper we investigate this relationship between efficiency and fairness in settings such as the above, and develop mechanisms that produce a good tradeoff between the two.

The first necessary step in analyzing the tradeoff between efficiency and fairness is to get a clear understanding of each of the notions. Efficiency is rather straightforward - it concerns the overall cost of obtaining the goal and is measured by totaling the costs of all team members. In the stochastic case - we take the expectation of this total.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

Fairness: Ex Ante and Ex Post.

Fairness concerns the *difference* in costs between team members - dissimilar costs are considered less fair while similar costs are considered fairer. Accordingly, we measure fairness as the ratio between the cost incurred by the agent incurring lowest cost and that incurred by the agent incurring the highest cost. In the stochastic case, it would seem that we can simply compare the expected costs, but this is not entirely true - as similar expectations need not result in expected similarity. Indeed, even if the expected costs of all team members are identical, their actual costs may frequently be very dissimilar. Furthermore, individuals often care more about the actual fairness than about fairness in expectation. For example, the common practice of splitting costs based on receipts is specifically aimed at achieving actual fairness. Thus, we distinguish between two notions of fairness: *ex-ante fairness* and *ex-post fairness*. Ex-ante fairness considers the difference in the agents' expected costs, whereas ex-post fairness considers the expected actual difference in costs.

The formal definition of the model, as well as the formal definitions of efficiency and fairness are provided in Section 2.

Contributions.

In this work we consider different possible policies[1] for reimbursement among the agents, and analyze their performance with respect to efficiency and fairness. Our goal is to develop policies that are both efficient and fair, to the extent possible.

First we consider the common practice of sharing costs based on receipts. We show that the inefficiency of this policy is unbounded. Furthermore, when taking the search costs into account, the unfairness of this policy is also unbounded. We thus present an alternative mechanism which is both fully efficient and fully ex-ante fair. It does not, however, provide ex-post fairness. Indeed, assuming that the actual cost of search cannot be tracked, full ex-post fairness cannot be obtained. We show, however, that ex-post unfairness can nonetheless be reduced. We present a family of policies that improve ex-post fairness, providing a tradeoff between ex-post fairness and efficiency. For this family we provide a closed form solution for the searching agent's optimal strategy, enabling the system designer to evaluate a wide range of alternative policies. Numerical illustrations are provided throughout the paper, illustrating the effect of the model and mechanisms' performance.

2. MODEL AND PERFORMANCE MEASURES

For exposition purposes we adopt the product purchase terminology. The same results hold for other costly search settings as well.

We consider a setting where a team $\mathcal{A} = \{A_1, ..., A_k\}$ of self interested, fully-rational agents have designated one of its members, call *the buyer*, to buy a specific product on behalf of the team. The agent can buy the product from any one of n stores. The price of the product at each store is assumed to be drawn at random from a common distribution, characterized by probability distribution function (p.d.f.) $f(x)$, with which all agents are familiar [33]. The actual price at each store is known only upon "visiting" the store (either physically or virtually), a process which is associated with some cost c (e.g. in gas, time, effort), termed *search cost*. It is assumed that at any time, the product can be purchased at any store visited so far, including that with the lowest price.

Reimbursement Policy.

Since the agent is operating on behalf of the entire team, it is reimbursed by the other group members, to some extent, according to some pre-defined *reimbursement policy*. The policy states how much money will be paid to the buyer by the other team members (who share this cost equally among themselves). In determining the reimbursement amount, the policy can take into account all available information, both a-priori and actual. We assume that the buyer can provide evidence, e.g. in the form a receipt, for the actual price paid for the product. Actual search costs, on the other hand, are assumed not be provable or otherwise available to the policy.

Strategy.

The price distribution and search cost, on the one hand, and the reimbursement policy, on the other, together determine the optimal strategy for the buyer. The buyer's strategy is characterized by a *stopping rule* which determines if to continue searching, and incur the associated costs, or stop searching and purchase at the lowest price found thus far. The optimal strategy for the agent is the one that minimizes its expected net cost, defined as the expenses incurred along the process (both for searching and for buying) minus the reimbursement received. Note, that the above characterization of the individual agent's problem extends the standard and widely used canonical costly search problem [10, 21, 19, 26, 29] to the setting with reimbursement.

Measures.

Obviously different reimbursement policies result in different individual search strategies and consequently different distributions of the total expense and its division among the group members. Therefore, for evaluating the different reimbursement policies we consider two measures of interest for the agents: efficiency and fairness.

Efficiency reflects the total net cost incurred by the team as a whole. Note that reimbursement is made internally between team members, so the cost of the team is fully captured by the buyer's total expenses for searching and buying (we use the term *expense* to denote out-of-pocket payment amounts, while *cost* denotes net-cost, i.e. expense minus reimbursement). Therefore, we define the efficiency of a given policy as the ratio between the minimal buyer's total expense, taken over all possible policies, and the expected total expense of the given policy. Formally, for a policy $\mathcal{P}$, let $expense^{\mathcal{P}}$(Buyer) be the expense of the buyer under this policy. Then, the efficiency of a policy $\mathcal{P}$ is defined as:

$$efficiency(\mathcal{P}) = \frac{\min_{\mathcal{Q}}\{E(expense^{\mathcal{Q}}(\text{Buyer}))\}}{E(expense^{\mathcal{P}}(\text{Buyer}))}$$

(where E is the expectation).

Fairness reflects the similarity in costs amongst the team members. Thus, fairness is measured as the ratio between the cost incurred by the agent incurring lowest cost and that incurred by the agent incurring the highest cost. As explained in the Introduction, we distinguish between two types of fairness: ex-ante fairness and ex-post fairness. Ex-ante fairness reflects the fairness amongst expected costs, while ex-post fairness reflects the expected fairness amongst actual costs. Formally, let $cost^{\mathcal{P}}(A_i)$ be the net cost of agent A_i under policy $\mathcal{P}$ (cost is taken to be 0 if the agent gains from the protocol). Then,

$$Fairness^{\text{ex-ante}}(\mathcal{P}) = \frac{\min_i\{E(cost^{\mathcal{P}}(A_i))\}}{\max_i\{E(cost^{\mathcal{P}}(A_i))\}}$$

and

$$Fairness^{\text{ex-post}}(\mathcal{P}) = E\left(\frac{\min_i\{cost^{\mathcal{P}}(A_i)\}}{\max_i\{cost^{\mathcal{P}}(A_i)\}}\right)$$

[1]We use the notions policy and mechanism interchangeably.

Note that both efficiency and fairness range between 0 and 1.

It is notable that a reimbursement rule that guarantees an ex-post fairness of 1 necessarily guarantees also an ex-ante fairness of 1. Nevertheless, for all other cases, the two are completely different. We define inefficiency and unfairness as reciprocal functions of efficiency and fairness, respectively.

The definition of efficiency has the minimal possible expense, taken over all possible policies, in the numerator. Luckily, there is one simple policy where this minimum is always obtained. Namely, the policy $\mathcal{P}_\emptyset$ in which the buyer is not reimbursed at all.

PROPOSITION 1. *The policy $\mathcal{P}_\emptyset$ guarantees the minimal possible expected buyer's total expense, i.e.*

$$E(\mathit{expense}^{\mathcal{P}_\emptyset}(\mathit{Buyer})) = \min_{\mathcal{Q}}\{E(\mathit{expense}^{\mathcal{Q}}(\mathit{Buyer}))\}$$

PROOF. By definition, the expense of the buyer under $\mathcal{P}_\emptyset$ is also its net cost. By assumption the buyer employs the strategy that minimizes its net cost. Thus, no strategy can obtain a lesser net-cost and hence also no lesser expense. □

Thus, to compute efficiency we use the buyer's expected expense under $\mathcal{P}_\emptyset$, denoted by $\mathit{expense}_\emptyset$, in the numerator.

3. RECEIPT SPLITTING POLICIES

The optimal strategy under $\mathcal{P}_\emptyset$ can be found in classic economic search theory [25]: the buyer should set a reservation value (i.e., a threshold) r, calculated as the solution for the following equation:

$$c = \int_{y=-\infty}^{r} (r-y)f(y)dy \tag{1}$$

The buyer should check prices in stores (in a random order, as they are all a priori alike) and terminate once running into a price lower than the reservation value r (or running out of stores to check). Intuitively, r is the value where the buyer is precisely indifferent: the expected marginal benefit from visiting another store exactly equals the cost incurred. It is notable that the decision rule is myopic, i.e., the value of r does not depend on the number of stores that can still be potentially explored [25].

In order to formulate the buyer's expected expense when using the above search strategy we first calculate the expected price at which the product is eventually purchased, denoted EV. This will also be the basis for the reimbursement policies developed later on. In order to calculate EV we first calculate the probability that the minimum price obtained throughout search, when using the reservation value r, is lesser than x, denoted $\bar{F}(x)$, calculated according to:

$$\bar{F}(x) = \begin{cases} 1-(1-F(x))^n & x > r \\ 1-\Big((1-F(r))^n + \frac{1-(1-F(r))^n}{F(r)}(F(r)-F(x))\Big) & x \le r \end{cases}$$

The case where $x > r$ requires that all n stores checked yield a price greater than x. The case $x \le r$ is calculated using the complementary probability, i.e., the probability that the best value obtained is greater than x. This corresponds to two possible scenarios. The first is where all n explored stores result in a price greater than r, i.e., with probability $(1-F(r))^n$. The second, is where the search terminates right after visiting the jth store, upon revealing a price y such that $x < y \le r$ (as otherwise, if $y > r$ the search should resume) and all the former $j-1$ stores checked returned a price greater than r (as otherwise the jth store is not reached). The probability for this latter case (for all values of $j \le n$) can be calculated using the geometric series $\prod_{j=1}^{n}(F(x)-F(r))F(r)^{j-1}$.

The first order derivative of $\bar{F}(x)$ is the probability distribution function of the minimum price, denoted $\bar{f}(x)$, and is given by $\bar{f}(x) = \frac{d(\bar{F}(x))}{dx}$. This enables the calculation of EV as:

$$EV = \int_{y=-\infty}^{\infty} y\bar{f}(y)dy \tag{2}$$

Using the above, the buyer's expected expense under $\mathcal{P}_\emptyset$, is given by:

$$\mathit{expense}_\emptyset = c \cdot \frac{1-(1-F(r))^n}{F(r)} + EV \tag{3}$$

where the first term is the expected cost incurred throughout the search, calculated as: $c\sum_{j=1}^{n}(1-F(r))^{j-1} = c\frac{1-(1-F(r))^n}{F(r)}$, as the number of stores checked is a geometric random variable bounded by n, with a success probability $F(r)$. The second term is the actual expense of buying the product.

When adding a reimbursement policy $\mathcal{P}$, there is no guarantee that the buyer's optimal search strategy is reservation-value based. Instead, it can take various forms (e.g., be based on several ranges of values for which additional stores are checked). Still, for some reimbursement policies we can prove the validness of the reservation-value based strategy structure. In particular, consider the common reimbursement policy, denoted $\mathcal{P}_{\mathit{split}}$, where the agents equally split the item's cost as indicated by the receipt, i.e., the buyer is reimbursed for $\frac{k-1}{k}$ of the cost of the item. The buyer's optimal strategy for this reimbursement is given in the following proposition.

PROPOSITION 2. *The buyer's optimal strategy under $\mathcal{P}_{\mathit{split}}$ is to use a reservation-value based strategy where the reservation value r is the solution to:*

$$c = \frac{1}{k} \cdot \int_{y=-\infty}^{r} (r-y)f(y)dy \tag{4}$$

PROOF. We first prove the reservation-value nature of the optimal strategy. Since recall is allowed then if the buyer prefers terminating search given the best known price v and n' remaining uncertain stores, it will also prefer terminating when the best known value is $v' < v$ (and n' remaining uncertain stores). The proof that the value of the reservation value does not depend on the number of remaining uncertain stores is inductive, showing that if with any number of remaining uncertain stores smaller than n' the optimal choice is to use reservation value r then so is the case with n' remaining uncertain stores. The reservation value when only one uncertain store is available derives from equating the search cost c with the expected improvement obtained by the additional search, i.e., the search resumes for any v for which $c < \frac{\int_{y=-\infty}^{v}(v-y)f(y)dy}{k}$, resulting in a reservation value r according to Equation 4. Now assume that the optimal reservation value to be used with any $n'' < n'$ uncertain stores is r and consider the agent's decision regarding checking one more store, if the best value it obtained so far is x and the number of uncertain stores is n''. If $x < r$ and the buyer checks one additional store, then regardless of the price obtained next the buyer will definitely terminate the search after the additional search (as it already has a price lower than r). This is equivalent to resuming the search when the best price obtained thus far is x and only one uncertain store is available. The latter choice however is not optimal according to the assumption that a reservation value r is used for any $n'' < n'$ uncertain available stores. Similarly, notice that the expected reduction in the agent's cost when resuming search with n' uncertain stores is greater than with $n'' < n'$ uncertain stores. Therefore, if resuming search is optimal for $x > r$, for n'', according to the inductive assumption, then so should be the optimal choice when with n'. $\square$

Note that a totally analogous proof actually shows that a similar behavior is obtained whenever the reimbursement is determined as a fraction of the item's actual cost. For any α, let $\mathcal{P}_\alpha$ be the policy where the buyer is reimbursed with an α fraction of the item's actual cost ($\mathcal{P}_{split} = \mathcal{P}_{\frac{k-1}{k}}$). The buyer's optimal strategy under $\mathcal{P}_\alpha$ is reservation based with a reservation value as in (4) with $\frac{1}{k}$ replaced by $1-\alpha$.

From (4) we obtain that, as expected, the reservation value used by the buyer increases when the number of agents in the group, k, increases. More generally, as the portion α that the buyer is being reimbursed increases, the buyer's reservation value increases. This is because the buyer's decision is based on comparing the cost of further search with the expected individual saving due to the increase in the receipt's amount. When being reimbursed a greater portion, the marginal savings from each improvement in purchase price decreases, hence the buyer is more reluctant to resume search. Furthermore, for $k = 1$ (or $\alpha = 0$) the agent's considerations are the same as with no reimbursement at all, and indeed Equation 4 is the same as (1).

Using the same analysis methodology given in former paragraphs, the expected expense of the buyer, under the $\mathcal{P} = \alpha$ rule, is given by:

$$E(cost^{\mathcal{P}_{split}}(\text{Buyer})) = c\frac{1-(1-F(r))^n}{F(r)} + (1-\alpha)\cdot EV \quad (5)$$

where r is calculated according to (4), replacing $1/k$ by $1-\alpha$. The first term in (5) is the expected cost incurred throughout the search carried out by the buyer. The second term is the buyer's share in the expense of buying the product (i.e., after reimbursed by the group), where EV is calculated according (2) using the appropriate reservation value. The expected reimbursement will be $\alpha \cdot EV$.

The above enables us to illustrate the problematic nature of the common "splitting-receipts" policy, and more generally $\mathcal{P}_\alpha$ policies. This approach results both in a non-efficient and unfair outcomes. Figure 1 depicts the three performance measures (efficiency, ex-ante fairness and ex-post fairness) as function of α for a setting of 5 stores, where product prices derive from the uniform distribution function over $(0, 1)$ and search cost is $c = 0.02$, for different number of agents k. In order to calculate the ex-post fairness, we first calculated the optimal buyer's strategy and then used simulation, averaging over 1 million results for each setting. As observed from Figure 1(a), efficiency decreases as the percentage of reimbursement increases. Furthermore, the maximum efficiency is achieved with no reimbursement, i.e., when $\alpha = 0$. Unlike with the efficiency, as the figure depicts, the maximum ex-ante fairness (graph (b)) and the maximum ex-post fairness (graph (c)) are achieved with some positive, however different than the natural $split$ (marked with dotted vertical lines), reimbursement. It is notable that in this specific case the number of agents in the team does not affect efficiency (unlike with the fairness measures), because the buyer is being reimbursed according to α rather than k. Also, as expected, the ex-post fairness does not reach 1, as this would mean an equal division of the expense in each instance.

Interestingly, we can prove that both the inefficiency and unfairness resulting from $\mathcal{P} = \alpha$-like reimbursement rules are unbounded.

PROPOSITION 3. *The inefficiency and unfairness in $\mathcal{P}_{split}$ are unbounded.*

PROOF. We begin with inefficiency. Consider the case where $f(x)$ is uniform between 0 and 1 (i.e., $f(x) = 1$, $\forall 0 \le x \le 1$, otherwise $f(x) = 0$) and the buyer is not bounded by the number of stores (i.e., $n \to \infty$). In this case the minimum buyer expense is achieved when the buyer uses $r' = \sqrt{2\cdot c}$, according to (1). Since the distribution of prices is uniform, we obtain $EV = r'/2$ (for any $r' \le 1$) and the expected number of stores visited is $1/r'$. Therefore $expense_\emptyset = r'/2 + c/r'$ which, after substituting $r' = \sqrt{2\cdot c}$ turns into $expense_\emptyset = \sqrt{2\cdot c}$. Similarly, under the $\mathcal{P}_{split}$ reimbursement policy the buyer's reservation value r'' is given by $r'' = \sqrt{2\cdot k\cdot c}$. Thus we obtain:

$$E(expense^{\mathcal{P}_{split}}(\text{Buyer})) = r''/2 + c/r'' = \frac{k\cdot c + c}{\sqrt{2\cdot k\cdot c}}$$

Therefore the inefficiency measure is: $\frac{\frac{k\cdot c+c}{\sqrt{2\cdot k\cdot c}}}{\sqrt{2\cdot c}} = \sqrt{k}/2 + 1/\sqrt{4k}$, hence we can increase the inefficiency indefinitely by increasing k.[2]

Next, we show that the ex-ante unfairness is unbounded, using the exact same setting as above. For a sufficiently large c value, we obtain $r' > 1$. In this case $EV = 1/2$, $E(cost^{\mathcal{P}_{split}}(\text{Buyer})) = 1/(2k) + c$ and $E(cost^{\mathcal{P}_{split}}(A_i)) = 1/(2k)\ \forall A_i \neq Buyer$. Therefore:

$$Fairness^{ex\text{-}ante}(\mathcal{P}_{split}) = \frac{1/(2k) + c}{1/(2k)} = 1 + 2kc$$

Therefore we can increase the ex-ante unfairness indefinitely by increasing k.

Since the ex-ante unfairness measure is unbounded, then necessarily the ex-post unfairness measure is unbounded too. $\square$

4. FULL EFFICIENCY POLICIES

Since the buyer can supply evidence only for the actual expense of purchase (i.e., the receipt amount), fully ex-post fairness can not be achieved. This is because, given the probabilistic nature of the search, each receipt amount can be associated with a wide range of possible accumulated search cost. Therefore, a reimbursement policy that relies only on the receipt amount as a decision parameter will necessarily fail to result in an equal cost for the buyer and the remaining of the agents for at least one possible outcome of accumulated search cost. Therefore, in this section we discuss reimbursement policies that guarantee ex-ante full fairness while keeping a fully efficient search. We first show two reimbursement policies that guarantee both full efficiency and full ex-ante fairness. Each policy influences differently the ex-post fairness. Hence, neither dominates the other. These two policies are then generalized into a single parametrized policy that corresponds to a class of policies aiming to balance costs in order to achieve high ex-post fairness.

The first policy, $\mathcal{P}_{fixed}(s)$, is characterized by the buyer being paid some fixed amount s, regardless of the outcome of the search process. In fact, this amount can be paid in advance, prior to carrying out the search process. The following proposition determines that the $\mathcal{P}_{fixed}(s)$ is efficient and determines the s which provides full ex-ante fairness.

PROPOSITION 4. *For $s = expense_\emptyset \cdot (k-1)/k$, the policy $\mathcal{P}_{fixed}(s)$ is fully efficient and fully ex-ante fair.*

PROOF. Since the buyer is payed a fixed amount, it accounts for whatever expense incurred beyond that amount. Therefore, its search optimization problem is identical to minimizing the total

[2]We can always find c small enough such that (1) will result in a reservation value $r'' < 1$, hence $EV = r''/2$.

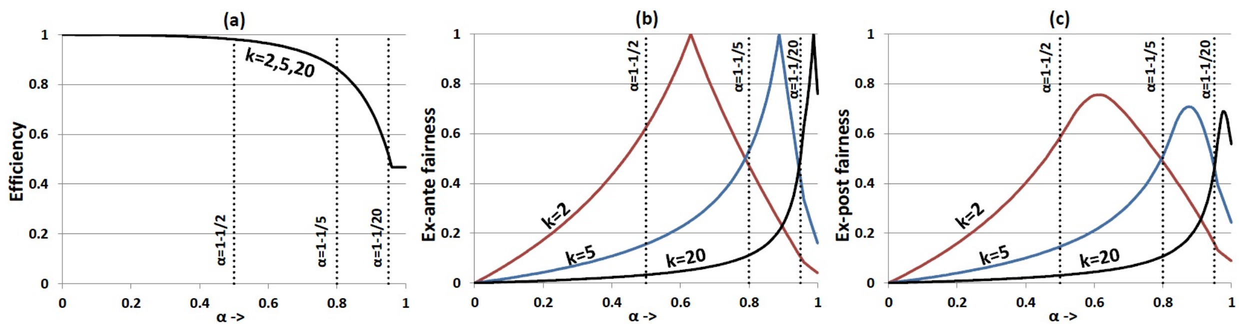

Figure 1: **The effect of α in the $\mathcal{P} = \alpha$ reimbursement policy on: (a) efficiency; (b) ex-ante fairness; and (c) ex-post fairness for different number of agents. The setting used is: $n = 5$, $c = 0.02$ and $f(x)$ uniformly distributed between 0 and 1.**

expense, i.e., similar to the one used with $\mathcal{P}_{\emptyset}$. Hence by Proposition 1, $\mathcal{P}_{\mathit{fixed}}(s)$ is fully efficient. It remains to determine that for $s = \mathit{expense}_{\emptyset} \cdot (k-1)/k$, the policy $\mathcal{P}_{\mathit{fixed}}(s)$ is fully ex-ante fair. This is straightforward — since the buyer uses the same strategy as with $\mathcal{P}_{\emptyset}$, its expected expense will be $\mathit{expense}_{\emptyset}$, and consequently its expected cost is $\mathit{expense}_{\emptyset}/k$, which equals the equal share of each of the other agents in the reimbursement. □

The policy $\mathcal{P}_{\mathit{fixed}}(s)$ is characterized by a fixed cost for the non-searching team members and a variable cost to the buyer (due to its variable search and purchase expenses). In fact, in this case, the uncertainty associated with the search is captured entirely by the buyer's cost, hence affecting ex-post fairness. In an attempt to obtain more balance between the two, we introduce a second policy, denoted $\mathcal{P}_{\mathit{bonus}}(s)$. Under $\mathcal{P}_{\mathit{bonus}}(s)$ the reimbursement is composed of two parts:

- Sharing: the buyer gets $(k-1)/k$ of actual amount paid for the good.
- Bonus: an amount s is added as a bonus if the actual price paid for the good is no more than the reservation price r, as determined in (1).

The following proposition shows that by choosing s appropriately we can guarantee full efficiency and fairness.

PROPOSITION 5. *Setting $s = \frac{c(k-1)}{F(r)k}$ provides that $\mathcal{P}_{\mathit{bonus}}(s)$ is fully efficient and fully ex-ante fair.*

PROOF. In order to prove efficiency, we prove that the buyer's strategy in this case is reservation value, and the optimal reservation value, denoted r', satisfies $r' = r$ according to (1). The proof is the same as the one given for Proposition 2, except that the expected benefit from resuming search, if the best known price is $v > r$, is given by $\frac{\int_{y=-\infty}^{r'}(r'-y)f(y)dy}{k} + s \cdot F(r)$, and $\frac{\int_{y=-\infty}^{r'}(r'-y)f(y)dy}{k}$ otherwise ($v \leq r$), resulting in:

$$c = \frac{\int_{y=-\infty}^{r'}(r'-y)f(y)dy}{k} + \frac{c(k-1)F(r)}{F(r)k}, v > r \quad (6)$$

When $v < r$, further search is necessarily not beneficial, as c is greater than $\frac{\int_{y=-\infty}^{r'}(r'-y)f(y)dy}{k}$ according to (1). Notice that (6) can be represented as:

$$ck = \int_{y=-\infty}^{r'}(r'-y)f(y)dy + c(k-1)$$

which after some mathematical manipulations becomes identical to Equation 1, hence the same reservation value will be used ($r' = r$).

Using the same methodology as in former section, the expected buyer's expense in this case is given by:

$$E(\mathit{cost}^{\mathcal{P}_{\mathit{bonus}}(s)}(\text{Buyer})) = c \cdot \frac{1-(1-F(r))^n}{F(r)} + \frac{\int_{y=-\infty}^{\infty} y\bar{f}(y)dy}{k} - s \cdot \bar{F}(r) = \frac{\mathit{expense}_{\emptyset}}{k}$$

and the expected cost of any of the remaining agents is:

$$E(\mathit{cost}^{\mathcal{P}_{\mathit{bonus}}(s)}(A_i)) = \frac{\int_{y=-\infty}^{\infty} y\bar{f}(y)dy}{k} + \frac{s \cdot \bar{F}(r)}{k-1} = \frac{\mathit{expense}_{\emptyset}}{k}$$

hence, the cost of all agents is equal.

□

One inherent shortcoming of the $\mathcal{P}_{\mathit{bonus}}(s)$ policy, from the ex-post fairness point of view, is that the buyer is being reimbursed for the search cost portion of its expense only in cases where it has actually managed to find a price below the reservation value according to (1). This suggests that in some cases it will be reimbursed only $(k-1)/k$ of the receipt amount, while in others it will be reimbursed substantially more than what it has actually spent (to compensate for all the times it has not been reimbursed its search cost). For example, consider the case of two agents and only one store that can be checked, where there is a probability of 0.5 for a price 0 and a probability 0.5 for a price 1000. The search cost is $c = 500$. In this example, in half of the cases the buyer's cost will be 1000 and the other agent's cost will be 500. In the remaining of the cases, the buyer's cost will be 0 and the other agent's cost will be 500. Both cases are highly unequal and thus highly unfair.

In order to assure a more balanced reimbursement for the costs of the searcher, we introduce a third policy, $\mathcal{P}_{\mathit{combined}}(s_1, s_2, \beta)$, according to which some portion of the reimbursement is fixed (as in the first policy), and some portion is given based on success (as in the second policy). Notice, however, that if some portion of the reimbursement provided upon "success" in the former policy, as given in Proposition 5, is traded for an initial independent reimbursement, then an efficient search cannot be guaranteed (the benefit from further search when obtaining a price r will not worth the cost of such search). Therefore, a greater incentive should be provided for the buyer to keep searching even when finding a price close to r. We manage to achieve this goal by reimbursing the agent less than the reimbursement used by $\mathcal{P}_{\mathit{bonus}}(\cdot)$ for the receipt amount. This way, the saving from further search, from the buyer's point of view, is greater. Specifically, under $\mathcal{P}_{\mathit{combined}}(s_1, s_2, \beta)$, the buyer is reimbursed:

- Fixed: a fixed amount of s_1.

- Sharing: the buyer gets a β fraction of actual amount paid for the good.
- Bonus: an amount s_2 is added as a bonus if the actual price paid for the good is no more than the reservation price r, as determined in (1).

The following proposition provides the proper choices of s_1 and s_2 (given β) which guarantee full efficiency and ex-ante fairness.

PROPOSITION 6. *Setting:*

$$s_1=\left(c\cdot\frac{1-(1-F(r))^n}{F(r)}+(1-\beta)\cdot\int_{y=-\infty}^{\infty}y\bar{f}(y)dy-s_2\cdot\bar{F}(r)\right)-\frac{expense_\emptyset}{k}$$

and

$$s_2=\frac{c\cdot\beta}{F(r)}$$

$\mathcal{P}_{combined}(s_1,s_2,\beta)$ *is fully efficient and fully ex-ante fair.*

PROOF. For the same considerations given in the proof for Proposition 5, the buyer will use a reservation-value rule where the optimal reservation value satisfies $c=(1-\beta)\cdot\int_{y=-\infty}^{r'}(r'-y)f(y)dy+s_2\cdot F(r)$, where r' is its cost-minimizing strategy. Substituting $s_2=\frac{c\cdot\beta}{F(r)}$, obtains:

$$c=(1-\beta)\cdot\int_{y=-\infty}^{r'}(r'-y)f(y)dy+c\cdot\beta$$

which after some mathematical manipulations becomes identical to Equation 1, hence the reservation value used by the buyer, r', satisfies $r'=r$ and achieves the maximum efficiency. In this case the expected cost of the buyer is:

$$E(cost^{\mathcal{P}_{combined}(s_1,s_2,\beta)}(Buyer))=c\cdot\frac{1-(1-F(r))^n}{F(r)} \quad (7)$$
$$+(1-\beta)\cdot\int_{y=-\infty}^{\infty}y\bar{f}(y)dy-s_2\cdot\bar{F}(r)-s_1$$

and after substituting s_1 in (7):

$$E(cost^{\mathcal{P}_{combined}(s_1,s_2,\beta)}(A_i))=expense_\emptyset/k$$

hence achieving full ex-ante fairness. $\square$

We note that the first two policies given in this section are actually specific cases of the third. The first uses $\beta=0$ and reimburse the agent regardless of "success". The second uses $\beta=(k-1)/k$ and reimburse the agent only upon "success".

Figure 2 depicts the results of ex-post fairness, achieved with the generalized policy described above, as a function of the value β used, for different number of search costs in a setting with two agents and 5 stores (i.e., $k=2$, $n=5$) (graph (a)) and for different number of stores in a setting with two agents and a search cost 0.01 (i.e., $k=2$, $c=0.01$) (graph (b)). The vertical dotted lines represent the two specific variants of the policy for which $\beta=0.5$ and $\beta=0$. In both cases prices of the product derive from the uniform distribution function over $(0,1)$. The figure was generated using a simulation of 1 million instances for each data point, as with Figure 1. As observed from the figure, while the specific case where $\beta=0.5$ yields for some settings the maximum ex-post efficiency, in others the maximum is achieved using a different β value. We note that the low efficiency achieved for relatively high and low β values is explained in this example in the fact that these are associated with a substantial pre-payments, either positive (for $\beta=0$) or negative (for $\beta=1$) hence in many cases (specific instances played) either of the agents incur a "negative" cost (meaning that the agent "gains" from the policy), in which case the fairness is by definition 0.

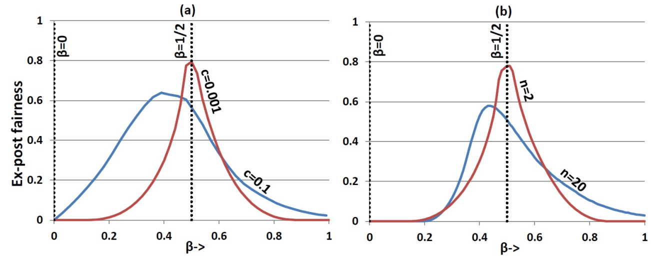

Figure 2: The effect of the percentage β out of the purchase price in the reimbursement of the non-searching agents on ex-post fairness for different: (a) search costs, where the setting used is: $n=5$, $k=2$; and (b) number of stores, where the setting used is: $c=0.01$, $k=2$. In both cases $f(x)$ is uniformly distributed between 0 and 1.

5. TRADING EFFICIENCY WITH FAIRNESS

A further improvement in the ex-post fairness can be achieved if one is willing to compromise on efficiency. In this section we present a family of policies that can improve ex-post fairness, but may degrade the efficiency, thus presenting a tradeoff between efficiency and ex-post fairness.

The main challenge in the design and analysis of reimbursement rules with no maximum efficiency guarantee is the derivation of the optimal search strategy for the buyer. With general reimbursement rules, the buyer's optimal strategy may be complex and hard to analyze. Luckily, the following Theorem establishes that for a large family of natural reimbursement policies the buyer's optimal search strategy is reservation-value based, and shows how this reservation value can be calculated.

THEOREM 1. *For a function $g:\mathbb{R}\rightarrow\mathbb{R}$, let $\mathcal{P}_{function}(g)$ be the reimbursement policy wherein the buyer is reimbursed a sum of $g(x)$, upon presenting a purchase receipt of x. Provided that $dg(x)/dx\leq 1$, the buyer's optimal search strategy is reservation value based, where the optimal reservation value r satisfies:*

$$c=\int_{y=-\infty}^{r}(r-g(r)-(y-g(y)))f(y)dy \quad (8)$$

PROOF. The proof generally resembles the one given for Proposition 2 thus we include only the differences. First, we prove the optimality of the reservation value strategy. The searcher's overall expense is divided into the expense it incurs for purchasing the product and the accumulated expense due to the search. Given the reimbursement policy $\mathcal{P}_{function}(g)$, the portion of the cost associated with paying the store becomes $x-g(x)$. Since $dg(x)/dx\leq 1$ then $d(x-g(x))/dx\geq 0$. Therefore, given a known price x, the searcher's benefit from improving the price to any value $z<x$ increases as x increases. Hence if it is optimal to resume search given that the best price is x then so is the optimal choice when the best price is $x'>x$. The remaining of the proof remains unchanged except for switching the benefit $\frac{\int_{y=-\infty}^{r}(r-y)f(y)dy}{k}$ with $\int_{y=-\infty}^{r'}(r'-g(r'))-(y-g(y)))f(y)dy$ whenever applicable. $\square$

Based on the above, we can formulate the expected cost of the buyer, when using the reservation value r:

$$E(cost^{\mathcal{P}}(Buyer))=c\frac{1-(1-F(r))^n}{F(r)}+\int_{y=-\infty}^{\infty}(y-g(y))\bar{f}(y)dy$$

The remaining agents will incur a cost of:

$$E(cost^{\mathcal{P}}(A_i))=\frac{1}{k}\cdot\int_{y=-\infty}^{\infty}g(y)\bar{f}(y)dy$$

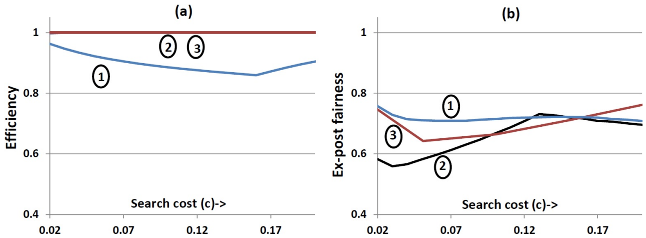

Figure 3: Trading Efficiency with Fairness: (a) efficiency; and (b) ex-post fairness; using non-efficiency-maximizing policies (marked as curves 1 and 2) and the best β-based efficiency-maximizing policy (curve 3). The setting used is: $n = 5, k = 2$ and $f(x) = 1, \forall 0 \leq x \leq 1$, otherwise $f(x) = 0$. The functions $g(x)$ used for the first two policies are given in the body of the paper.

Theorem 1 allows us to define and analyze a large set of reimbursement policies - including ones with less than full efficiency, and choose the one exhibiting the most preferred performance tradeoffs. It is interesting to note that all the policies presented in the previous section ($\mathcal{P}_{fixed}(\cdot), \mathcal{P}_{bonus}(\cdot)$ and $\mathcal{P}_{combined}(\cdot,\cdot,\cdot)$) are special cases of the $\mathcal{P}_{function}(g)$ family.

Figure 3 depicts performance (efficiency and ex-post fairness) exhibited by with three $\mathcal{P}_{function}(g)$ policies:

1. $\mathcal{P}_{function}(g_1)$ with $g_1(x) = (0.6 + 0.5 \cdot c) \cdot x$.
2. $\mathcal{P}_{function}(g_2)$ with
$$g_2(x) = \begin{cases} e^{-4\cdot x} & \text{x>r} \\ \frac{(x+c\cdot(1-(1-r)^n)/r)\cdot(k-1)}{3.8} & otherwise \end{cases}$$
where r is the reservation value calculated according to (1).
3. $\mathcal{P}_{function}(g_3)$ which is $\mathcal{P}_{combined}(s_1, s_2, \beta)$ such that for each cost c we take the β value that results in the maximum ex-post fairness for this class of policies and s_1 and s_2 are computed according to Proposition 6.

As can be observed from the figure, the new reimbursement rules improve the ex-post fairness for a substantial portion of the settings (graph (b)). This however comes on the expense of the efficiency measure (graph (a)), therefore the system designer should consider the tradeoff between the two and choose accordingly.

6. RELATED WORK

Historically, efficiency has been the main objective of teamwork and multi-agent systems [34]. Recently, as agents frequently represent human individuals, the importance of *fairness*, as an independent goal, has gained recognition within the MAS literature (see [13, 9], and in particular the review given in [14]). Many works use the “cake cutting” setting as a model for considering fairness in multi-agent systems (see [27] for a recent survey). The cake cutting model, introduced in [32], postulates a continuously-divisible good (a.k.a. “the cake”) to be divided among a group of agents. Different agents may place different values on the different pieces of the cake, and the goal is to divide the cake among the agents in a fair way - under some suitable definition of fairness. For the cake cutting model, the tension between fairness and efficiency has been studied through the notion of the “price of fairness”, which is defined as the ratio between the maximal possible social welfare if no fairness is required and maximal possible welfare when fairness is also required (this is analogous to our notion of inefficiency). It has been established that the price of fairness in cake cutting can be, at times, unbounded, depending on the exact model and the fairness criteria [6, 1]. Accordingly, an important line of research has been in devising algorithms that provide fairness, while optimizing welfare [11, 4]. Another line of research has considered the problem of devising *truthful* cake cutting algorithms, in cases where agents may misrepresent their true valuation functions [8]. The cake cutting literature assumes, for the most part, a non-transferable utility setting, wherein agent cannot pay each other.

The cake cutting model assumes a continuously-divisible good. Fair allocations of indivisible goods has also been considered, providing both algorithms and hardness results [3, 23]. Other works consider mediated negotiation procedures that support negotiating agents in reaching Pareto efficient and fair agreements, e.g., in bilateral multi-issue negotiation [22] and computational models which allow agents to find the most desirable solution according to certain definitions of fairness or optimality [15].

While the definition of fairness in the above body of literature is mostly similar to the goal our research attempts to achieve, our work studies the problem of a team search, where fairness does not depend on the way resources are allocated, but rather on the amount of effort an agent will put in the search process. The analysis we provide relies heavily on understanding the optimal search strategy that will be used by an agent under different reimbursement policies, and the resulting effect over the fairness achieved. To the best of our knowledge, no prior work has considered this important problem in that sense.

Costly search of the kind used in this paper is a prevalent theme in MAS [18, 19, 30, 21]. The idea is that agents need to consume some of their resources to disambiguate the uncertainty associated with the different alternatives and options available to them. In some sense, the basic sequential search model can be seen as part of the field of planning under uncertainty, hence related to Markov decision processes (MDP) and decentralized Markov decision processes [2], as the goal is to maximize the expected cumulative reward, which is also the objective in costly search. However, the analysis provided by “search theory” using threshold-based solutions, whenever proved to be optimal, is simpler and can be derived with a substantially lesser complexity compared to solving MDPs. Alongside models of a single agent search, several models of group or team search have been introduced [28]. Most of the work in this area has focused on a representative agent, operating on behalf of the team [31, 17, 24, 7, 5]. These however assumed that the designated agent is fully cooperative, and as such focused in maximizing efficiency, i.e., attempted to extract an optimal search strategy. As such, fairness was not a consideration in these works at all.

Fairness is a major topic of interest in the economics literature, which is out of the scope of this paper to review (see [12, 16, 20]). In the economic-search literature, however, we are not aware of consideration of fairness.

7. DISCUSSION

In this work we considered efficiency and fairness in team search, and the possible tradeoffs between the two. We distinguished between two separate notions of fairness - ex-ante and ex-post. The first measures the similarity between the expected costs, while the second measures the expected similarity in actual costs. We believe this is an important distinction, as humans frequently care more about ex-post fairness - while ex-ante fairness is algorithmically easier to obtain.

We consider the case where a single buyer needs to purchase a good on behalf of the entire team, and analyze the efficiency and fairness resulting from different reimbursement policies. We show that the common policy of splitting costs based on purchase receipts

- primarily aimed at ex-post fairness - may result in severe degradation in efficiency. We thus present two alternative families of reimbursement policies. The one family guarantees full efficiency and full ex-ante fairness, but may lack in ex-post fairness. This family uses a mix of fixed sum and receipt sharing reimbursement, together with a possible bonus if a certain price level is achieved. By choosing the right parameters for each of these three types of reimbursement we can strike the proper balance - incentivising the buyer to search for a good price, while also maintaining fairness. The second family of policies allows more complex reimbursement policies, including variable percentage sharing of the purchase price and multiple bonus levels. With this type of policies, we can further improve ex-post fairness, at times - at the expense of efficiency. Thus, we allow the MAS designer to trade efficiency for fairness, as appropriate for each setting.

This work considered a concrete MAS setting where fairness is a frequent requirement in practice. We demonstrated the substantial impact fairness can have on efficiency, and the intricate tradeoffs and relationships between the two. We believe that such tradeoffs may also arise in many other MAS settings, and should be acknowledged and analyzed. On a more general level, we believe that as MAS systems become more integrated with human interaction, *fairness* should become a more central and driving notion in the planning of such systems. Humans are very sensitive to fairness, and failing to take it into account may result in poor performance of such systems in practice.

Acknowledgments

This work was partially supported by ISF grant 1083/13, the Israeli Ministry of Industry and Trade under project RESCUE.

8. REFERENCES

[1] Y. Aumann and Y. Dombb. The efficiency of fair division with connected pieces. In *Proc. WINE*, pages 26–37, 2010.

[2] D. Bernstein, D. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of markov decision processes. *MOR*, 27(4):819–840, 2002.

[3] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *JAIR*, 32:525–564, 2008.

[4] S. J. Brams, M. Feldman, J. K. Lai, J. Morgenstern, and A. D. Procaccia. On maxsum fair cake divisions. In *Proc. AAAI*, pages 1285–1291, 2012.

[5] K. Burdett and D. A. Malueg. The theory of search for several goods. *JET*, 24(3):362–376, 1981.

[6] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. The efficiency of fair division. *TOCS*, 50(4):589–610, 2012.

[7] J. A. Carlson and R. P. McAfee. Joint search for several goods. *JET*, 32(2):337–345, 1984.

[8] Y. Chen, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, pages 284–297, 2012.

[9] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J. Padget, S. Phelps, J. A. Rodriguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1):3–31, 2006.

[10] M. Chhabra, S. Das, and D. Sarne. Expert-mediated search. In *Proc. AAMAS*, pages 415–422, 2011.

[11] Y. J. Cohler, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Optimal envy-free cake cutting. In *Proc. AAAI*, pages 626–631, 2011.

[12] A. Dannenberg, T. Riechmann, B. Sturm, and C. Vogt. Inequity aversion and individual behavior in public good games: An experimental investigation. Technical Report 07-034, 2007.

[13] S. de Jong and K. Tuyls. Human-inspired computational fairness. *JAAMAS*, 22(1):103–126, 2011.

[14] S. de Jong, K. Tuyls, and K. Verbeeck. Fairness in multi-agent systems. *Knowledge Eng. Review*, 23(2):153–180, 2008.

[15] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proc. AAMAS*, pages 177–184, 2003.

[16] E. Fehr and K. M. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868, 1999.

[17] J. Gatti. Multi-commodity consumer search. *JET*, 86(2):219–244, 1999.

[18] A. Grosfeld-Nir, D. Sarne, and I. Spiegler. Modeling the search for the least costly opportunity. *EJOR*, 197(2):667–674, 2009.

[19] N. Hazon, Y. Aumann, S. Kraus, and D. Sarne. Physical search problems with probabilistic knowledge. *Artificial Intelligence*, 196:26–52, 2013.

[20] D. Kahneman, J. L. Knetsch, and R. H. Thaler. Fairness and the assumptions of economics. *The Journal of Business*, 59(4):285–300, 1986.

[21] J. Kephart and A. Greenwald. Shopbot economics. *JAAMAS*, 5(3):255–287, 2002.

[22] M. Li, Q. B. Vo, and R. Kowalczyk. Searching for fair joint gains in agent-based negotiation. In *Proc. AAMAS*, pages 1049–1056, 2009.

[23] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proc. EC*, pages 125–131, 2004.

[24] E. Manisterski, D. Sarne, and S. Kraus. Enhancing cooperative search with concurrent interactions. *JAIR*, 32(1):1–36, 2008.

[25] J. McMillan and M. Rothschild. Search. In *Proc. Handbook of Game Theory with Economic Applications*, pages 905–927, 1994.

[26] P. Morgan and R. Manning. Optimal search. *Econometrica*, 53(4):923–944, 1985.

[27] A. D. Procaccia. Cake cutting: not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.

[28] I. Rochlin and D. Sarne. Information sharing under costly communication in joint exploration. In *Proc. AAAI*, pages 847–853, 2013.

[29] M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4):689–711, 1974.

[30] D. Sarne. Competitive shopbots-mediated markets. *ACM Trans. Econ. Comput.*, 1(3):17:1–17:41, 2013.

[31] D. Sarne, E. Manisterski, and S. Kraus. Multi-goal economic search using dynamic search structures. *JAAMAS*, 21(2):204–236, 2010.

[32] H. Steinhaus. Sur la division pragmatique. *Econometrica*, 17:315–319, 1949.

[33] R. Waldeck. Search and price competition. *Journal of Economic Behavior and Organization*, 66(2):347–357, 2008.

[34] G. Weiss, editor. *Multiagent systems: a modern approach to distributed artificial intelligence*. MIT Press, 1999.