JAAMAS manuscript No.
(will be inserted by the editor)

Efficiency and Fairness in Team Search with Self-Interested Agents

Igor Rochlin · Yonatan Aumann · David Sarne · Luba Golosman

Received: date / Accepted: date

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.

Keywords Multi-Agent Exploration · Joint Exploration · Cooperation · Fairness in MAS · Teamwork

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

I. Rochlin
School of Computer Science, College of Management, Rishon LeZion 7502501, Israel.
E-mail: igor.rochlin@gmail.com

Y. Aumann
Department of Computer Science, Bar-Ilan University, Ramat-Gan 5290002, Israel.
E-mail: aumann@cs.biu.ac.il

D. Sarne
Department of Computer Science, Bar-Ilan University, Ramat-Gan 5290002, Israel.
E-mail: sarned@cs.biu.ac.il

L. Golosman
Department of Computer Science, Bar-Ilan University, Ramat-Gan 5290002, Israel.
E-mail: golosml@macs.biu.ac.il

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 among 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 - frequently accept that others benefit from their work, they also expect that the distribution of costs among the team members be "fair", in the sense that costs incurred are, more or less, evenly distributed among the team 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 whom 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. 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 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 pay in full for her assigned book, then they may end up paying different amounts - which could be considered 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.

Fairness concerns the *difference* in costs between team members - dissimilar costs are considered less fair than similar costs.

Accordingly, we measure fairness as the ratio between the cost incurred by the agent with the lowest cost and that incurred by the agent with 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 splitting 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 alternative mechanisms which are both fully efficient and fully ex-ante fair. They do 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.

2 Model and Performance Measures

For exposition purposes we adopt the product purchase terminology, which is the most natural domain for our model.

We consider a setting where a team $\mathcal{A} = \{A_1, ..., A_k\}$ of self-interested, fully-rational agents has designated one of its members, called *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 [51,52,25]. 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 resources, time, effort), termed *search cost*. It is assumed that, at any time, the product can be purchased at any store visited so far, including the one with the lowest price (this is termed *recall* [9,31]).

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 of a receipt, for the actual price paid for the product. Actual search costs, on the other hand, are assumed to be unprovable or otherwise unavailable to the policy.

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

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 whether to continue searching, and incur the associated costs, or to 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 [11, 27, 24, 32, 46] to the setting with reimbursement.

Measures Obviously different reimbursement policies result in different individual search strategies and consequently different distributions on the total expense and its division among the group members. For evaluating the different reimbursement policies we consider two measures of interest: 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 buyer's minimal total expense, taken over all possible policies, and the expected total expense of the given policy. Formally, for a policy $\mathcal{P}$, let $\mathit{expense}^{\mathcal{P}}(\text{Buyer})$ be the expense of the buyer under this policy. Then, the efficiency of a policy $\mathcal{P}$ is defined as:

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

(where E is the expectation).

The definition of efficiency has the minimal possible expense, taken over all possible policies, in the numerator. Fortunately, 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 buyer's minimal possible expected total expense, that is,*

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

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.

Fairness reflects the similarity in costs amongst the team members. Thus, fairness is measured as the ratio between the cost incurred by the agent with the lowest cost and that incurred by the agent with 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 among expected costs, while ex-post fairness reflects the expected fairness among actual costs. Formally, let $\mathit{cost}^{\mathcal{P}}(A_i)$ be the net cost of agent A_i under policy $\mathcal{P}$. Then,

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

and

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

Note that the term in the denominator is always positive (as otherwise the aggregated expense incurred is negative, which is impossible). It is possible, however, that the nominator may turn out to be negative under some policy P. This corresponds to a situation where one of the agents ends up with a negative (either expected or actual) expense and this can happen only to the buyer (as this is the only agent that gets paid by others). Such a scenario where the buyer incurs no cost whereas the other agents incur in total more than the total expense for buying the product, while quite rare, captures the most possible unfairness. Therefore, we take the net cost of the buyer to be 0 if it gains from the protocol, so all cases of the latter type will obtain zero fairness according to the ex-ante and ex-post fairness measures. Therefore 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 different. We define inefficiency and unfairness as the reciprocal functions of efficiency and fairness, respectively.

3 Efficient (Optimal) Strategy

The optimal strategy under $\mathcal{P}_\emptyset$ can be found in classic economic search theory [31,21,28, 54]. Here, a searcher is facing n opportunities, where each opportunity (or, store) is associated with a value v which is a priori unknown to it — only the distribution $f(y)$ from which values are drawn is known. The searcher's goal is to maximize its expected revenue (or, depending on the application, minimize its expected expense) while the true value of any opportunity can be revealed for a cost c. In our model, the buyer's goal is to minimize its overall expected expense (the price paid to buy the product plus to costs to use the platform and query the sellers).

The solution for the problem in its expected-cost-minimization form the buyer sets a reservation value (i.e., a threshold) $r_\emptyset$, calculated as the solution to the following equation:

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

The buyer should check prices in stores (in an arbitrary order, as they are all a priori alike) and terminate once running into a price lower than the reservation value $r_\emptyset$ (or running out of stores to check). Intuitively, $r_\emptyset$ 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 $r_\emptyset$ does not depend on the number of stores that can still be potentially explored [31]. The proof for the stationarity of $r_\emptyset$ (based on [54], however substantially shortened based on the fact the stores available to the buyer are associated with common probability distribution function and search cost) is given in the following paragraph.

Since recall is allowed then if the buyer prefers terminating search given the best known value v and n' remaining uncertain stores it will also prefer that choice 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 greater than n' the optimal choice is to use reservation value $r_\emptyset$ 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 of that store with the expected improvement obtained by the additional search, i.e., the search resumes for any v for which $c < \int_{y=-\infty}^{v}(v-y)f_i(y)dy$, resulting in a reservation value $r_\emptyset$ according to Equation 1. Now assume that the optimal reservation value to be used with any $n'' > n'$ uncertain stores is $r_\emptyset$ and consider the buyer's decision regarding cheking one more store, if the best value it obtained so far is x and the number of uncertain available store is n''. If $x < r_\emptyset$ and the buyer executes one additional search, then regardless of the value obtained next the buyer will definitely terminate the search after the additional search (as it already has a value greater than $r_\emptyset$). This is equivalent to resuming the search when the best value obtained thus far is x and only one uncertain opportunity is available. The latter choice however is not optimal according to the assumption that a reservation value $r_\emptyset$ is used for any $n'' > n'$ uncertain available stores. Similarly, if $x > r_\emptyset$ and the buyer chooses to terminate the search, then terminating is dominated by exploring one more store and then terminating.

In order to compute the buyer's expected expense when using the above 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 the search, when using the reservation value $r_\emptyset$, is at most x, denoted $\bar{F}(x)$, calculated according to:

$$\bar{F}(x) = \begin{cases} 1-(1-F(x))^n & x > r_\emptyset \\ 1-\left((1-F(r_\emptyset))^n + \frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)}(F(r_\emptyset)-F(x))\right) & x \le r_\emptyset \end{cases} \quad (2)$$

The case $x > r_\emptyset$ is calculated using the complementary probability, i.e., the probability that no value lower than x was found. This requires that all n stores are checked and yield a price greater than x. Similarly, the case $x \le r_\emptyset$ is calculated based on the two possible scenarios according to which no value lower than x is found. The first is where all n explored stores result in a price greater than $r_\emptyset$, i.e., with probability $(1-F(r_\emptyset))^n$. The second is where the search terminates right after visiting the jth store, upon revealing a price y such that $x < y \le r_\emptyset$ (as otherwise, if $y > r_\emptyset$ the search should resume) and all the former $j-1$ stores checked returned a price greater than $r_\emptyset$ (as otherwise the jth store is not reached). The probability for this latter case (for all values of $j \le n$) is the geometric series $\sum_{j=1}^{n}(F(x)-F(r_\emptyset))F(r_\emptyset)^{j-1}$ which equals $\frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)}(F(r_\emptyset)-F(x))$.

The first derivative of $\bar{F}(x)$ is the probability distribution function of the minimum price, which we denote $\bar{f}(x)$. Thus, we obtain that EV, the expected minimum price, is:

$$EV = \int_{y=-\infty}^{\infty} y\bar{f}(y)dy \quad (3)$$

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

$$\textit{expense}_\emptyset = c \cdot \frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)} + EV \quad (4)$$

where the first term is the expected cost incurred throughout the search, calculated as: $c\sum_{j=1}^{n}(1-F(r_\emptyset))^{j-1} = c\frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)}$, the chance of visiting the jth store is $(1-F(r_\emptyset))^{j-1}$. The second term is the actual price of buying the product.

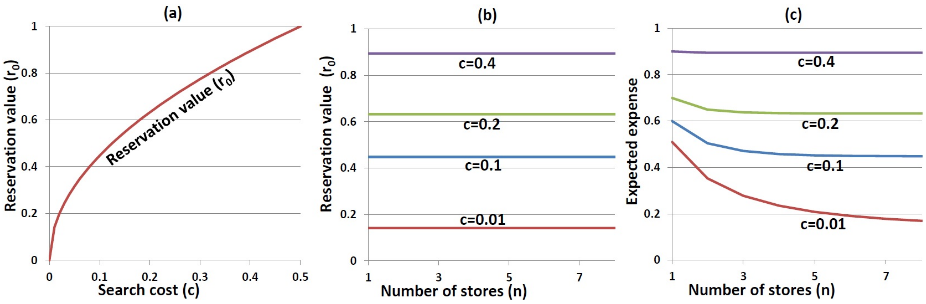

Fig. 1 (a) Reservation value as a function of the search cost; (b) Reservation value as a function of the number of stores, for different search costs; (c) Expected expense as a function of the number of stores, for different search costs.

Figure 1(a) depicts the expected-expense-minimizing reservation value $r_\emptyset$ of a buyer as a function of the search cost c. As expected, the value of $r_\emptyset$ increases as c increases (see Equation 1). It is notable that the value of $r_\emptyset$ does not depend on the number of stores available to the buyer, n, as explained above. This latter characteristic is illustrated in Figure 1(b), which depicts the expected-expense-minimizing reservation value as a function of the number of stores available to the buyer, for different search costs. The value of n does, however, influence the buyer's expected expense (see Figure 1(c)), such that the buyer's expected expense decreases as the number of stores increases (for a fixed c value), since it is less likely to run into situations where it is optimal to resume search yet no additional stores are available. The decrease in the expected expense becomes more moderate as the value of c increases, since the resulting increase in the reservation value used by the buyer suggests a shorter search process, on average, thus the constraint on the search extent becomes less influential.

4 Receipt Splitting Policy

When adding a reimbursement policy $\mathcal{P}$, there is no guarantee that the buyer's optimal search strategy is reservation-value based. For some reimbursement policies, however, the strategy is reservation-value based as we now show. Consider the common reimbursement policy, denoted $\mathcal{P}_{split}$, where the agents equally split the product's cost as indicated by the receipt, i.e., the buyer is reimbursed for $\frac{k-1}{k}$ of the cost of the product. The buyer's optimal strategy for this reimbursement policy is given in the following proposition.

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

$$c = \frac{1}{k} \cdot \int_{y=-\infty}^{r} (r-y)f(y)dy \tag{5}$$

Proof We first prove the reservation-value nature of the optimal strategy for the buyer. Then we continue with an inductive proof, showing that if the reservation value calculated according to (5) is indeed the optimal strategy for any number of remaining stores $m' > m$ then this reservation value should also be used when the number of remaining stores is m.

Since recall is allowed along the search process, the buyer's strategy is a mapping $S(\mathbf{X}, m) \to \{resume, terminate\}$, where $\mathbf{X}$ is the set of prices obtained so far and m

is the number of potential remaining uncertain stores. Since the buyer is interested merely in the minimum store price, its strategy is affected only by the minimum value in $\mathbf{X}$, hence the strategy can be defined in the form $S(v, m) \rightarrow \{terminate, resume\}$, where v is the minimum value in $\mathbf{X}$. Obviously, if according to the optimal strategy the buyer needs to resume the search with the pair (v, m), then the same should be true for any (v', m) with $v' > v$. Similarly, if according to the optimal strategy the search should terminate at pair (v, m), then the same should hold for any (v'', m) with $v'' < v$. Therefore, for each given number of remaining uncertain stores m, the optimal search strategy of the buyer can be characterized by the reservation value r^m such that the buyer should resume the search if the best price obtained so far is above r^m and otherwise it should terminate the search.

We now prove by induction on m, that $r^m = r$ (of (5)), beginning with the case of $m = 1$. If the best price obtained so far by the buyer is v, then exploring the one last store will incur search cost c and the expected expense of the buyer (excluding the search cost incurred so far) will be:

$$\frac{\int_{y=-\infty}^{\infty} \min(v, y) f(y) dy}{k} + c$$

(where y is the price obtained from the additional store). On the other hand, if the search is terminated, the expected expense of the buyer is given by:

$$\frac{v}{k} \tag{6}$$

Therefore, the buyer should visit the one last store if and only if:

$$\frac{v}{k} > \frac{\int_{y=-\infty}^{\infty} \min(v, y) f(y) dy}{k} + c$$

Substituting $v \cdot \int_{y=-\infty}^{\infty} f(y) dy$ for v (as $\int_{y=-\infty}^{\infty} f(y) dy = 1$), obtains:

$$\frac{\int_{y=-\infty}^{\infty} v \cdot f(y) dy}{k} > \frac{\int_{y=-\infty}^{\infty} \min(v, y) f(y) dy}{k} + c$$

which transforms to:

$$\frac{\int_{y=-\infty}^{\infty} (v - \min(v, y)) f(y) dy}{k} - c > 0 \tag{7}$$

Since the left-hand-side of (7) is an increasing function of v, the buyer should explore one last store whenever the value of v is greater than the value of r that satisfies (5). This completes the proof for $m = 1$.

Now assume the same r (according to (5)) holds for any $m' > m$, for some m, and consider the buyer's decision regarding exploring one more store, if the best price it obtained so far is v and the number of stores that can still be potentially explored is m. If $v < r$ and the buyer executes one additional search, then regardless of the price obtained next the buyer will definitely terminate the search after the additional search (as it already has a value lower than r and according to the induction assumption the optimal strategy thereafter is the reservation value r). Therefore the expense obtained from one additional search is given by:

$$\frac{\int_{y=-\infty}^{\infty} (v - \min(v, y)) f(y) dy}{k} - c$$

Alas, since the latter term increases as v increases, and obtains zero for $v = r$ (according to (5)), then for $v < r$ the term obtains a negative value, hence an additional search cannot be the preferred choice.

Similarly, consider the case where $v > r$ for m and the buyer chooses to terminate the search process. We show that terminating is dominated by exploring one more store and then terminating. This results immediately from $m = 1$, as we known that for $m = 1$ and $v > r$ the marginal change in the expected expense, if visiting one store (and then necessarily terminating) is greater than the cost c. Therefore, the optimal strategy for m is also a reservation value strategy and the optimal reservation value is calculated, once again, according to 5. □

Similar behavior is obtained whenever the reimbursement is determined as any other fraction of the product's actual cost. This is, for any α, let $\mathcal{P}_\alpha$ be the policy where the buyer is reimbursed with an α fraction of the product's actual cost (in particular, $\mathcal{P}_{split} = \mathcal{P}_{\alpha=\frac{k-1}{k}}$). The buyer's optimal strategy under $\mathcal{P}_\alpha$ is given in the following theorem.

Theorem 1 *The buyer's optimal strategy under $\mathcal{P}_\alpha$ is to use a reservation-value based strategy where the reservation value r is the solution to:*

$$c = (1-\alpha) \cdot \int_{y=-\infty}^{r} (r-y) f(y) dy \tag{8}$$

Proof The proof is identical to the one given for Proposition 2, replacing $\frac{1}{k}$ by $1-\alpha$ throughout. □

From (5) we obtain that, as expected, the reservation value used by the buyer increases when the number of agents in the group, k, increases, i.e., it will invest less effort in getting a good deal whenever splitting the payment for the searched product with more others. More generally, from (8), 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 its individual cost of further search with the expected individual saving from a lesser price. When being reimbursed a greater portion, the marginal saving from each improvement in price decreases, hence the buyer is more reluctant to resume search. For $k = 1$ (or $\alpha = 0$) the agent's considerations are the same as with no reimbursement at all, and indeed Equation 8 is the same as (1). Interestingly, in the specific case of $\mathcal{P}_{split}$ that is given by $\alpha = \frac{k-1}{k}$, the greater the size of the team (k) the lesser the efficiency achieved. This is because $\frac{k-1}{k}$ increases as k increases, hence in order to keep the equality in (8), we need to further increase r. Therefore, the buyer becomes reluctant to search further, even in scenarios where the current known price is substantial.

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

$$E\left(cost^{\mathcal{P}_\alpha}(\text{Buyer})\right) = c\frac{1-(1-F(r))^n}{F(r)} + (1-\alpha) \cdot EV \tag{9}$$

where r is calculated according to (8). The first term in (9) is the expected cost incurred throughout the search. The second term is the buyer's share in the product's price (i.e., after reimbursed by the group), where EV is calculated according to (3) using the appropriate reservation value.

The above analysis enables us to illustrate the problematic nature of the common "splitting-receipts" policy, and more generally $\mathcal{P}_\alpha$ policies. These policies result in both non-efficient and unfair outcomes. Figure 2 depicts the three performance measures (efficiency, ex-ante fairness and ex-post fairness) as a function of α for a setting of 5 stores, where product prices derive from different distribution functions (according to row) and the search cost is $c = 0.02$, for different numbers of agents, k. In order to calculate the ex-post fairness, we first calculated the buyer's optimal strategy and then used simulations to generate actual expenses, averaging the resulting ex-post fairness over 1 million runs for each setting. This procedure of calculating the ex-post fairness based on 1 million runs was used also in all the other graphs in the paper that depict that measure. As observed in Figure 2 (left column graphs), efficiency decreases as the percentage of reimbursement increases. Starting from some very large α value, the buyer's portion in the product's actual cost becomes negligible compared to the search cost incurred (i.e., the search cost becomes the dominant factor in the buyer's expected expense, according to Equation 9). In such cases, the buyer would rather terminate the search process upon checking the first store, regardless of the value obtained.[2] Therefore once α reaches the threshold at which the buyer limits its search to a single store, the efficiency of the search process remains constant for any greater α value. We note that the curve is the same for any value of k, since the search is executed by a single agent (the buyer) and with the $\mathcal{P}_\alpha$ policy the other agents' influence on its strategy is fully captured by the value of α (unlike with the fairness measure). Furthermore, we observe that the maximum efficiency is achieved with no reimbursement, i.e., when $\alpha = 0$, as expected. The maximum ex-ante fairness (middle column graphs) and the maximum ex-post fairness (right column graphs) are achieved with some positive reimbursement, however different than the natural $split$ (marked with dotted vertical lines, each corresponding to a $\mathcal{P}_{split}$ with a different number of agents). In fact, the maximum fairness is always achieved with an α value satisfying $\alpha > \frac{k-1}{k}$, since otherwise the buyer is not even equally compensated for the cost of the product itself. Also, as expected, the ex-post fairness does not reach 1, as this would mean an equal division of the expense in each instance. Finally, we note that as observed from the figures, the maximal ex-ante and ex-post fairness are achieved at very similar values of α, regardless of k for all three distributions checked. We discuss this in more detail later in this section. We emphasize that all the above behaviors discussed for the ex-post fairness rely on averages (of 1 million runs) hence the actual ex-post in an instance is likely to be highly influenced by the variance. Indeed, for the million runs we checked (for each α value checked) the standard deviation (as a percentage of the average value depicted in the figure) is substantial: ranging from 27% to 120% for the uniform distribution case, 38% to 72% for the exponential case and 2% to 8% for the normal case.

The fact that the $\mathcal{P}_\alpha$ reimbursement policy does not guarantee maximum efficiency calls for an analysis of the tradeoff between efficiency and fairness. Presumably, one might be willing to compromise the expected efficiency if a greater fairness can be achieved. While we discuss such tradeoffs in detail in Section 6, we provide a similar illustration for the $\mathcal{P}_\alpha$ policy as illustrated in Figure 3. The figure depicts the maximum efficiency one can possibly obtain (i.e., over all possible α value) when using the $\mathcal{P}_\alpha$ reimbursement policy as a function of the ex-ante fairness (graph (a)); and the ex-post fairness (graph (b)) for a setting of 5 stores and search cost is $c = 0.02$, for different number of agents. The product prices derive from the uniform distribution over $(0, 1)$ throughout, and the standard deviation (as a percentage of the depicted average value) ranging from 27% to 120% in graph (b). As can be observed from the figure, given that all other setting parameters equal, an increase in the number of agents k results in an increase in the maximum achievable efficiency for any requested level

[2] To illustrate, consider the case where $\alpha \to 1$. Here, the expected expense is fully attributed to the search cost, hence search should be constrained to the minimum possible (one store).

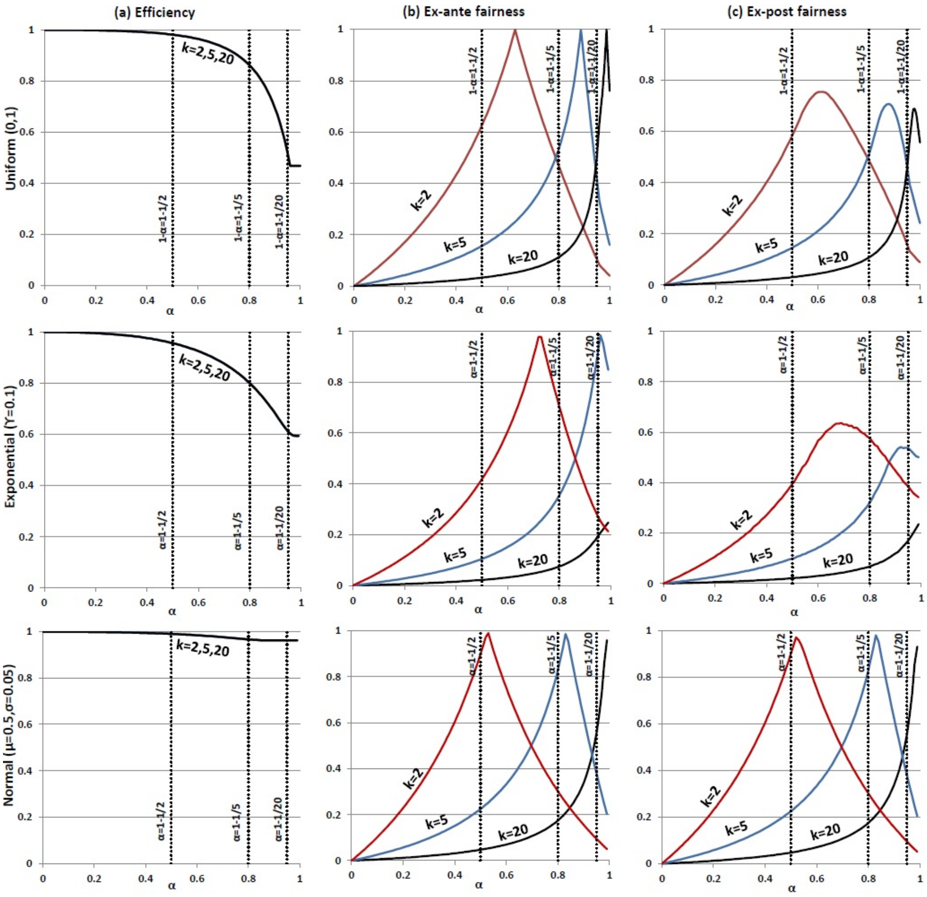

Fig. 2 The effect of α in the $\mathcal{P}_\alpha$ reimbursement policy on: (a) efficiency; (b) ex-ante fairness; and (c) ex-post fairness for different numbers of agents. The setting used is: $n = 5$, $c = 0.02$ and the distribution function differs between rows.

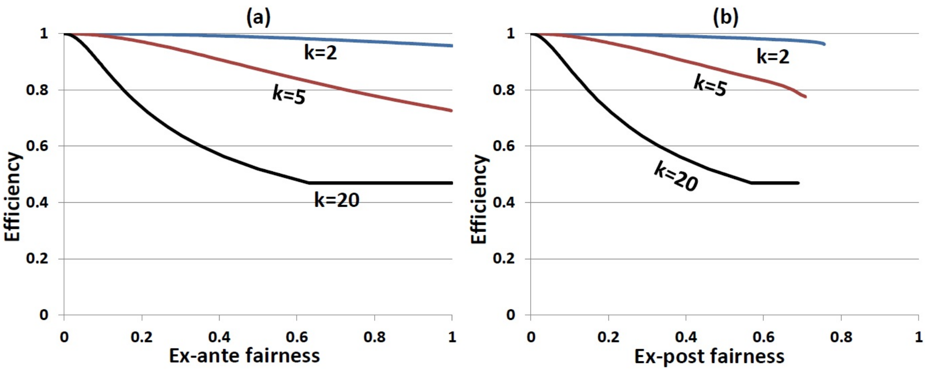

Fig. 3 The efficiency in the $\mathcal{P}_\alpha$ reimbursement policy as function: (a) ex-ante fairness; and (b) ex-post fairness for different number of agents. The setting used is: $n = 5$, $c = 0.02$ and $f(x)$ uniformly distributed between 0 and 1.

of ex-ante and ex-post fairness. We note that the maximum achievable ex-post fairness in this setting is less than 0.8 (depending on the number of agents) hence the graphs in part (b) of the figure terminate towards that value. Overall, we see that the maximum achievable efficiency decreases as the required ex-post increases. This enables choosing the appropriate α value to be used, based on the preferences between efficiency and fairness.

All in all, we see that the common $\mathcal{P}_{split}$ policy is far from being optimal. This is highlighted by the following proposition.

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 and the number of stores is infinity (i.e., $n = \infty$). In this case, by (1), the minimum buyer expense is achieved when the buyer uses reservation value $r_\emptyset$ with:

$$c = \int_{y=-\infty}^{r_\emptyset} (r_\emptyset - y) f(y) dy$$

and since $f(x)$ is uniform between 0 and 1:

$$c = \int_{y=0}^{r_\emptyset} (r_\emptyset - y) dy = r_\emptyset^2 - \frac{r_\emptyset^2}{2}$$

Thus:

$$r_\emptyset = \sqrt{2 \cdot c} \tag{10}$$

Since the distribution of prices is uniform, we obtain $EV = r_\emptyset/2$ (for any $r_\emptyset \leq 1$) and the expected number of stores visited is $1/r_\emptyset$. Therefore $expense_\emptyset$ is given by:

$$expense_\emptyset = r_\emptyset/2 + c/r_\emptyset$$

which, after substituting $r_\emptyset = \sqrt{2 \cdot c}$, turns into:

$$expense_\emptyset = r_\emptyset/2 + c/r_\emptyset = \frac{(\sqrt{2 \cdot c})^2 + 2c}{2\sqrt{2 \cdot c}} = \sqrt{2 \cdot c}$$

Similarly, under the $\mathcal{P}_{split}$ reimbursement policy the buyer's reservation value r_{split}, according to (5), is given by:

$$c = \frac{\int_{y=-\infty}^{r_{split}} (r_{split} - y) f(y) dy}{k}$$

and since $f(x)$ is uniform between 0 and 1:

$$r_{split} = \sqrt{2 \cdot k \cdot c} \tag{11}$$

Since we can always find c small enough such that (11) will result in a reservation value $r_{split} < 1$, we obtain $EV = r_{split}/2$. Thus, the expected expense is given by:

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

	Distribution		
Number of agents	Uniform	Normal	Exponential
2	0.992667	0.999976	0.987531
5	0.998044	0.999988	0.998692
20	0.999964	1	1

Table 1 Spearman's rank correlation coefficient of the ex-ante and ex-post fairness.

Therefore the inefficiency measure is:

$$\frac{1}{\textit{efficiency}(\mathcal{P}_{\textit{split}})} = \frac{\frac{k\cdot c+c}{\sqrt{2\cdot k\cdot c}}}{\sqrt{2\cdot c}} = \sqrt{k}/2 + 1/\sqrt{4k},$$

hence we can increase the inefficiency indefinitely by increasing k.

Next, we show that the ex-ante unfairness is also unbounded, using the exact same setting as above. For a sufficiently large c value, we obtain $r_{\emptyset} > 1$. In this case $EV = 1/2$,

$$E\left(\textit{cost}^{\mathcal{P}_{\textit{split}}}(\text{Buyer})\right) = 1/(2k) + c$$

and for $A_i \neq$ Buyer:

$$E\left(\textit{cost}^{\mathcal{P}_{\textit{split}}}(A_i)\right) = 1/(2k).$$

Therefore the unfairness measure is:

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

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

Since the ex-ante unfairness is unbounded, then necessarily the ex-post unfairness is also unbounded. □

We further studied the similarity between the behavior of the ex-ante and the ex-post fairness in $\mathcal{P}_\alpha$ policy using Spearman's rank correlation coefficient. For this purpose we used the settings that were used for Figure 2. The rank correlation coefficient of two variables measures the extent to which, as one variable increases, the other variable tends to increase, without requiring that the increase be represented by a linear relationship. Spearman's rank correlation coefficient rank is depicted in Table 1. For each distribution, the rank correlation coefficient increases with the number of agents and is very close to 1 for all price distributions and numbers of agents. Recall that the ex-ante fairness can be efficiently derived using a formula, while the ex-post fairness can only be obtained using extensive and time-consuming simulations. The high rank correlation coefficients show that indeed, it is safe to use the ex-ante formula as a prediction of the most fair setting of the policy for the ex-post fairness. Therefore, using this phenomena, a system designer can determine the value of α that maximizes the ex-post-fairness using the calculated ex-ante fairness without any simulation.

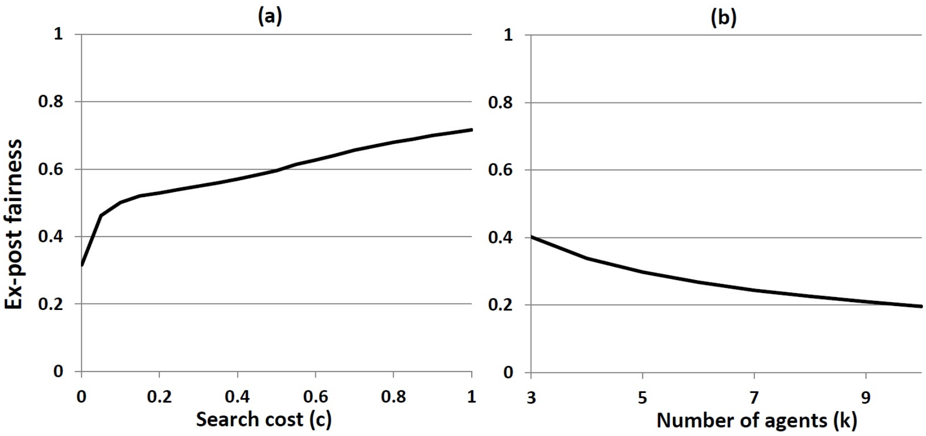

Fig. 4 The ex-post fairness in the $\mathcal{P}_{\mathit{fixed}}(\hat{s})$ reimbursement policy as a function of: (a) the search cost, where the setting used is: $n = 5$, $k = 2$; and (b) the number of agents, where the setting used is: $n = 5$, $c = 0.1$. In both cases $f(x)$ is uniformly distributed between 0 and 1.

5 Optimal Ex-ante Fairness and Optimal Efficiency Policies

In this section we present reimbursement policies that guarantee full ex-ante fairness and complete efficient. We first introduce two reimbursement policies obtaining this goal. Each policy influences the ex-post fairness differently. 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 better ex-post fairness.

The first policy, $\mathcal{P}_{\mathit{fixed}}(s)$, 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 $\mathcal{P}_{\mathit{fixed}}(s)$ is efficient and determines the s which provides full ex-ante fairness.

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

Proof Since the buyer is payed a fixed amount, it incurs all expenses beyond this amount. Therefore, its search optimization problem is identical to minimizing the total expense, i.e., identical to the one used with $\mathcal{P}_{\emptyset}$. Hence, by Proposition 1, $\mathcal{P}_{\mathit{fixed}}(s)$ is fully efficient (for any s). 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:

$$E\left(\mathit{cost}^{\mathcal{P}_{\mathit{fixed}}(\hat{s})}(\text{Buyer})\right) = \mathit{expense}_{\emptyset} - \mathit{expense}_{\emptyset} \cdot \left(\frac{k-1}{k}\right) = \frac{\mathit{expense}_{\emptyset}}{k},$$

which equals the equal share of each of the other agents in the reimbursement:

$$E\left(\mathit{cost}^{\mathcal{P}_{\mathit{fixed}}(\hat{s})}(A_i \neq \text{Buyer})\right) = \frac{\mathit{expense}_{\emptyset} \cdot \left(\frac{k-1}{k}\right)}{k-1} = \frac{\mathit{expense}_{\emptyset}}{k}$$

Therefore, the policy $\mathcal{P}_{\mathit{fixed}}(\hat{s})$ is fully ex-ante fair. $\square$

Figure 4 depicts the ex-post fairness of the $\mathcal{P}_{\mathit{fixed}}(\hat{s})$ reimbursement policy as a function of the search cost for a setting of 5 stores and 2 agents (graph (a)), and the number of agents for a setting of 5 stores with a search cost of $c = 0.1$ (graph (b)). The product prices derive from the uniform distribution over $(0, 1)$ throughout, and the standard deviation (as a percentage of the depicted average value) decreases from 96% to 23% in graph (a) and increases from 79% to 129% in graph (b). As observed from graph (a), the increase in search cost results in an increase in ex-post fairness. This behavior is explained by the fact that an increase in the search cost results in an increase in the reservation value and hence a decrease in the expected number of checked stores. Therefore, the variance in the search expenses of the buyer decreases and consequently the ex-post fairness increases. With regards to the number of agents (graph (b)), this number does not influence the calculation of the expected expense $\mathit{expense}_{\emptyset}$. On the other hand, an increase in the number of agents results in an increase in the fixed payment $\hat{s}$ to the buyer (according to Proposition 4) and at the same time a decrease in the expense of each non-buying agent A_i. Thus, the fixed expenses of all agents decrease, while the variable expense of the buyer remains unchanged. Thus, ex-post fairness decreases.

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

– Sharing: the buyer gets $(k-1)/k$ of the actual amount paid for the product.
– Bonus: an amount b is added if the actual price paid for the product is no more than the reservation price $r_{\emptyset}$, as determined in (1).

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

Proposition 5 *Setting $\hat{b} = \frac{c(k-1)}{F(r_{\emptyset})\cdot k}$ provides that $\mathcal{P}_{\mathit{bonus}}(\hat{b})$ 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 based, and the optimal reservation value, denoted r_{bonus}, satisfies $r_{bonus} = r_{\emptyset}$ according to (1). The proof that the buyer's strategy is reservation value is the same as in the proof given for Proposition 2. Hence we only need to prove that the same reservation value is used for any number of remaining stores and that this reservation value is calculated according to (1). The proof of this part follows the methodology used in the proof of Proposition 2, this time, however, taking into consideration the expected bonus-based reimbursement.

Beginning with the case of $m = 1$, the expected cost (excluding the search cost incurred so far) from resuming the search, in case the best price so far is $v \geq r_{\emptyset}$, is given by:

$$\frac{\int_{y=-\infty}^{\infty} \min(v, y) f(y) dy}{k} - \hat{b} \cdot F(r_{\emptyset}) + c$$

On the other hand, if the search is terminated, the expected expense of the buyer is given by:

$$\frac{v}{k} \tag{12}$$

Therefore, the buyer should visit the one last store if and only if:

$$\frac{v}{k} > \frac{\int_{y=-\infty}^{\infty} \min(v, y) f(y) dy}{k} - \hat{b} \cdot F(r_{\emptyset}) + c$$

Substituting $v \cdot \int_{y=-\infty}^{\infty} f(y)dy$ for v (as $\int_{y=-\infty}^{\infty} f(y)dy = 1$), obtains:

$$\frac{\int_{y=-\infty}^{\infty} v \cdot f(y)dy}{k} > \frac{\int_{y=-\infty}^{\infty} \min(v,y)f(y)dy}{k} - \hat{b} \cdot F(r_{\emptyset}) + c$$

which transforms to:

$$\frac{\int_{y=-\infty}^{v} (v-y)f(y)dy}{k} + \hat{b} \cdot F(r_{\emptyset}) > c \tag{13}$$

Substituting $\hat{b} = \frac{c(k-1)}{F(r_{\emptyset}) \cdot k}$ results in:

$$\frac{\int_{y=-\infty}^{v} (v-y)f(y)dy}{k} + \frac{c(k-1)F(r_{\emptyset})}{F(r_{\emptyset})k} > c \tag{14}$$

Thus (14) can be represented as:

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

or

$$\int_{y=-\infty}^{v} (v-y)f(y)dy > c$$

which is identical to Equation 1.

When the best price so far is $v < r_{\emptyset}$, further search is necessarily not beneficial, as:

$$\int_{y=-\infty}^{v} (v-y)f(y)dy < \int_{y=-\infty}^{r_{\emptyset}} (r_{\emptyset}-y)f(y)dy = c$$

according to (1). Hence the reservation value r_{bonus} for the case of $m = 1$ will be calculated according to (1), i.e., $r_{bonus} = r_{\emptyset}$. From this point, we can apply the same inductive proof used for Proposition 2 to prove that the same r_{bonus} holds for any number of remaining stores $m > 1$. Therefore, by Proposition 1, $\mathcal{P}_{bonus}(\hat{b})$ is fully efficient.

In order to prove that $\mathcal{P}_{bonus}(\hat{b})$ is fully ex-ante fair, we follow the methodology used in the proof of Proposition 4. The expected buyer's expense in this case is given by:

$$E\left(cost^{\mathcal{P}_{bonus}(\hat{b})}(\text{Buyer})\right) = c \cdot \frac{1-(1-F(r_{\emptyset}))^n}{F(r_{\emptyset})} + \frac{\int_{y=-\infty}^{\infty} y\bar{f}(y)dy}{k} - \hat{b} \cdot \bar{F}(r_{\emptyset})$$

Now, substituting $b = \frac{c(k-1)}{F(r_{\emptyset}) \cdot k}$ and $\bar{F}(r_{\emptyset}) = 1-(1-F(r_{\emptyset}))^n$ (according to (2)) obtains:

$$E\left(cost^{\mathcal{P}_{bonus}(\hat{b})}(\text{Buyer})\right) = c \cdot \frac{1-(1-F(r_{\emptyset}))^n}{F(r_{\emptyset})} + \frac{\int_{y=-\infty}^{\infty} y\bar{f}(y)dy}{k} - \\ - \frac{(k-1)}{k} \cdot c \cdot \frac{1-(1-F(r_{\emptyset}))^n}{F(r_{\emptyset})} = \frac{expense_{\emptyset}}{k}$$

where the last equality is based on the fact that $expense_{\emptyset} = c \cdot \frac{1-(1-F(r_{\emptyset}))^n}{F(r_{\emptyset})} + \int_{y=-\infty}^{\infty} y\bar{f}(y)dy$ (according to (4)). The expected expense of any of the remaining agents is:

$$E\left(cost^{\mathcal{P}_{bonus}(\hat{b})}(A_i)\right) = \frac{\int_{y=-\infty}^{\infty} y\bar{f}(y)dy}{k} + \frac{\hat{b} \cdot \bar{F}(r_{\emptyset})}{k-1}$$

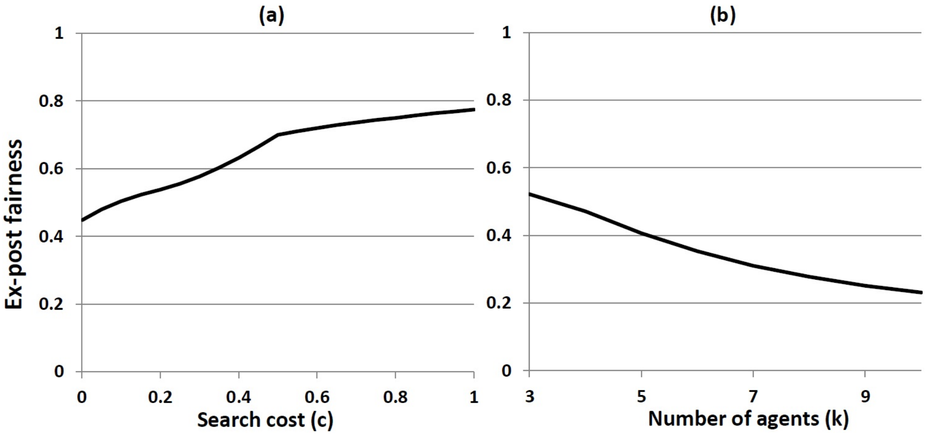

Fig. 5 The ex-post fairness in the $\mathcal{P}_{bonus}(\hat{b})$ reimbursement policy as a function: (a) the search cost, where the setting used is: $n = 5$, $k = 2$; and (b) number of agents, where the setting used is: $n = 5$, $c = 0.1$. In both cases $f(x)$ is uniformly distributed between 0 and 1.

and, once again, after substituting $\bar{F}(r_\emptyset) = 1 - (1 - F(r_\emptyset))^n$ (based on (2)) and $b = \frac{c(k-1)}{F(r_\emptyset)\cdot k}$, we obtain:

$$E\left(cost^{\mathcal{P}_{bonus}(\hat{b})}(A_i)\right) = \frac{\int_{y=-\infty}^{\infty} y\bar{f}(y)dy}{k} + \frac{c\cdot(1-(1-F(r_\emptyset))^n)}{F(r_\emptyset)\cdot k} = \frac{expense_\emptyset}{k}.$$

Hence, the expected expense of all agents is equal. □

Figure 5 depicts the ex-post fairness in the $\mathcal{P}_{bonus}(\hat{b})$ reimbursement policy as a function of the search cost (graph (a)) and the number of agents (graph (b)) (for the same settings used in Figure 4). The product prices derive from the uniform distribution over $(0, 1)$ throughout, and the standard deviation (as a percentage of the depicted average value) decreases from 61% to 11% in graph (a) and increases from 48% to 81% in graph (b). As observed from the figure, the ex-post fairness increases as the search cost increases, and decreases as the number of agents increases. The explanations of these phenomena are similar to those given when analyzing Figure 4.

One inherent shortcoming of the $\mathcal{P}_{bonus}(b)$ policy 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. This means that in some cases the buyer 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 for 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 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 ex-post unfair.

In order to achieve a more balanced reimbursement for the costs of the buyer, we introduce a third policy, $\mathcal{P}_{combined}(s, b, \beta)$, 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). Specifically, under $\mathcal{P}_{combined}(s, b, \beta)$, the buyer is reimbursed:

- Fixed: a fixed amount of s.
- Sharing: the buyer gets a β fraction of the actual amount paid for the product.
- Bonus: an amount b is added if the actual price paid for the good is no more than the reservation price $r_\emptyset$, as determined in (1).

The following proposition provides the proper choices of s and b (given β) that guarantee full efficiency and ex-ante fairness.

Proposition 6 *For every β, setting:*

$$\hat{s} = c \cdot \frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)} + (1-\beta) \cdot \int_{y=-\infty}^{\infty} y\bar{f}(y)dy - b \cdot \bar{F}(r_\emptyset) - \frac{\mathit{expense}_\emptyset}{k} \tag{15}$$

and

$$\hat{b} = \frac{c \cdot \beta}{F(r_\emptyset)} \tag{16}$$

the policy $\mathcal{P}_{combined}(\hat{s}, \hat{b}, \beta)$ is fully efficient and fully ex-ante fair.

Proof This proof follows the proof given for Proposition 5 except that the calculation of the expected expense now takes into consideration the additional reimbursement components, therefore we only include the differences. In order to prove efficiency, we prove that the buyer's strategy in this case is reservation value based, and the optimal reservation value, denoted $r_{combined}$, satisfies $r_{combined} = r_\emptyset$ according to (1). The expected cost (excluding the search cost incurred so far) from resuming the search, in case the best price so far is $v \geq r_\emptyset$, is given by:

$$(1-\beta) \int_{y=-\infty}^{\infty} \min(v,y) f(y)dy - \hat{b} \cdot F(r_\emptyset) - \hat{s} + c$$

On the other hand, if the search is terminated, the expected expense of the buyer is given by:

$$(1-\beta) \cdot v - \hat{s}$$

Therefore, the buyer should visit the one last store if and only if:

$$(1-\beta) \cdot v - \hat{s} > (1-\beta) \int_{y=-\infty}^{\infty} \min(v,y) f(y)dy - \hat{b} \cdot F(r_\emptyset) - \hat{s} + c$$

which transforms to:

$$(1-\beta) \int_{y=-\infty}^{v} (v-y) f(y)dy + \hat{b} \cdot F(r_\emptyset) > c \tag{17}$$

Substituting $\hat{b} = \frac{c \cdot \beta}{F(r_\emptyset)}$ results in:

$$(1-\beta) \int_{y=-\infty}^{v} (v-y) f(y)dy + \frac{c \cdot \beta \cdot F(r_\emptyset)}{F(r_\emptyset)} > c \tag{18}$$

which after some basic mathematical manipulations transforms into:

$$\int_{y=-\infty}^{v} (v-y) f(y)dy > c$$

which is identical to Equation 1.

When the best price so far is $v < r_\emptyset$, further search is necessarily not beneficial, as:

$$\int_{y=-\infty}^{v} (v-y)f(y)dy < \int_{y=-\infty}^{r_\emptyset} (r_\emptyset - y)f(y)dy = c$$

according to (1). Hence the reservation value used by the buyer, $r_{combined}$, satisfies $r_{combined} = r_\emptyset$ and maximum efficiency is achieved.

In order to prove full ex-ante fairness, we formulate the expected cost of the buyer:

$$E\left(cost^{\mathcal{P}_{combined}(\hat{s},\hat{b},\beta)}(\text{Buyer})\right) = c \cdot \frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)} + (1-\beta) \cdot \int_{y=-\infty}^{\infty} y\bar{f}(y)dy - \hat{b} \cdot \bar{F}(r_\emptyset) - \hat{s} \tag{19}$$

where the first term is the expected cost incurred throughout the search, calculated as:

$$c\sum_{j=1}^{n}(1-F(r_\emptyset))^{j-1} = c\frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)},$$

since the chance of visiting the jth store is $(1-F(r_\emptyset))^{j-1}$. The second term is the $(1-\beta)$ fraction of the actual price paid for the product by buyer. The third term is the bonus $\hat{b}$ received if the actual price paid is not greater than the reservation price $r_\emptyset$. The last term is the fixed amount $\hat{s}$ that is paid to the buyer regardless of the outcome of the search process.

After substituting $\hat{s}$, according to (15), in (19), we obtain:

$$E\left(cost^{\mathcal{P}_{combined}(\hat{s},\hat{b},\beta)}(\text{Buyer})\right) = \frac{expense_\emptyset}{k}.$$

The expected expense of any of the remaining agents is:

$$E\left(cost^{\mathcal{P}_{combined}(\hat{s},\hat{b},\beta)}(A_i \neq \text{Buyer})\right) = \frac{\beta \cdot \int_{y=-\infty}^{\infty} y\bar{f}(y)dy}{k-1} + \frac{\hat{b} \cdot \bar{F}(r_\emptyset)}{k-1} + \frac{\hat{s}}{k-1} \tag{20}$$

which after substituting $\hat{s}$, according to (15), becomes:

$$E\left(cost^{\mathcal{P}_{combined}(\hat{s},\hat{b},\beta)}(A_i \neq \text{Buyer})\right) = \frac{1}{k-1} \cdot \left(\int_{y=-\infty}^{\infty} y\bar{f}(y)dy + c \cdot \frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)} - \frac{expense_\emptyset}{k}\right)$$

However, from (3) and (4) we known that $expense_\emptyset = \int_{y=-\infty}^{\infty} y\bar{f}(y)dy + c \cdot \frac{1-(1-F(r_\emptyset))^n}{F(r_\emptyset)}$, therefore:

$$E\left(cost^{\mathcal{P}_{combined}(\hat{s},\hat{b},\beta)}(A_i \neq \text{Buyer})\right) = \frac{expense_\emptyset - \frac{expense_\emptyset}{k}}{k-1} = \frac{expense_\emptyset}{k}$$

hence achieving full ex-ante fairness. □

We note that the first two policies given in this section are specific cases of the third. The first uses $\beta = 0$ and reimburses the agent regardless of "success". The second uses $\beta = (k-1)/k$ and reimburses the agent only upon "success".[3]

[3] In this case $\hat{b} = \frac{c \cdot \beta}{F(r_\emptyset)} = \frac{c\cdot(k-1)}{k \cdot F(r_\emptyset)}$ which is equal to the bonus specified in Proposition 5 and the substitution of $\beta = \frac{k-1}{k}$ and $\hat{b} = \frac{c\cdot(k-1)}{k\cdot F(r_\emptyset)}$ results in $\hat{s} = 0$.

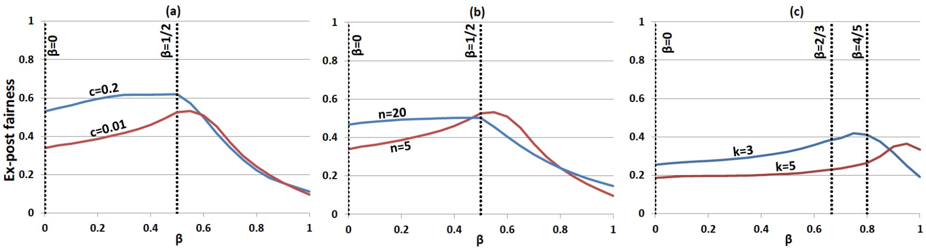

Fig. 6 The effect of the percentage β out of the purchase price in the reimbursement of the non-searching agents on ex-post fairness for different: (a) search costs, where the setting used is: $n = 5$, $k = 2$; (b) number of stores, where the setting used is: $c = 0.01$, $k = 2$; and (c) number of agents, where the setting used is: $c = 0.01$, $n = 5$. In all three cases $f(x)$ is uniformly distributed between 0 and 1.

Figure 6 depicts the results of ex-post fairness achieved with the generalized policy described above, as a function of the value β used, for several different search costs c in a setting with two agents and 5 stores (graph (a)); for several different numbers of stores n in a setting with two agents and a search cost of 0.01 (graph (b)); and for several different numbers of agents k, in a setting with 5 stores and a search cost $c = 0.01$ (graph (c)). The vertical dotted lines represent the two specific variants of the policy for which $\beta = \frac{k-1}{k}$, i.e., the pure bonus equivalent policy, and $\beta = 0$, i.e., the pure fixed payment policy. In all three graphs, prices of the product derive from the uniform distribution function over $(0, 1)$, and the standard deviation (as a percentage of the depicted average value) ranging from 13% to 118% in graph (a), from 55% to 118% in graph (b), and from 73% to 115% in graph (c). As observed from the figure, while the specific case where $\beta = \frac{k-1}{k}$ yields for some settings the maximum ex-post efficiency, in others, the maximum is achieved using different β values. We note that the low efficiency achieved for relatively high and low β values is explained in this example by the fact that these are associated with substantial pre-payments, either positive (for $\beta = 0$) or negative (for $\beta = 1$), hence in many cases (specific instances played) one of the agents incurs a "negative" cost (meaning that the agent "gains" from the policy), in which case the fairness is by definition 0.

6 Trading Efficiency for Fairness

Since the buyer can supply evidence only for the actual expense of purchase (i.e., the receipt amount), fully ex-post fairness cannot be achieved. This is because, given the probabilistic nature of the search, each receipt amount can be associated with a 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 identical cost for the buyer and the remaining of the agents for at least one possible outcome of accumulated search cost.

An 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 first challenge in the design and analysis of reimbursement rules with no maximum efficiency guarantee is the determination of the optimal search strategy for the buyer. With general reimbursement rules, the buyer's optimal strategy may be complex. 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 2 *For a function $g : \mathbb{R} \rightarrow \mathbb{R}$, let $\mathcal{P}_{function}(g)$ be the reimbursement policy wherein the buyer is reimbursed $g(x)$ upon presenting a purchase receipt of x. Provided that $dg(x)/dx \leq 1$, i.e., the lower the price found the greater the absolute reimbursement received, the buyer's optimal search strategy is reservation value based, wherein the optimal reservation value r satisfies:*

$$c = \int_{y=-\infty}^{r} (r - g(r) - (y - g(y)))f(y)dy \tag{21}$$

Proof This proof generally resembles the one given for Proposition 2. Thus we include only the differences. First, we prove the optimality of a reservation value strategy for this case. 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 purchasing the product becomes $x - g(x)$. Since $dg(x)/dx \leq 1$ then $d(x - g(x))/dx \geq 0$. Therefore, given that the best price so far is $\hat{x}$, the searcher's benefit from improving the price to any value $x < \hat{x}$ increases as $\hat{x}$ increases. Hence if it is optimal to resume the search given that the best price so far is $\hat{x}$, then it is also optimal when the best price is any $x > \hat{x}$. Thus, the optimal strategy is reservation value based.

The remainder of the proof remains unchanged except for substituting 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. □

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

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

where the first term is the expected cost incurred by the search and the second term is the net payment for the product (i.e., excluding the reimbursement received).

Any remaining agent A_i will incur a cost of:

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

Theorem 2 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 of 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:

(a) for $\mathcal{P}_{fixed}(\hat{s})$, $g(y) = \hat{s}$.

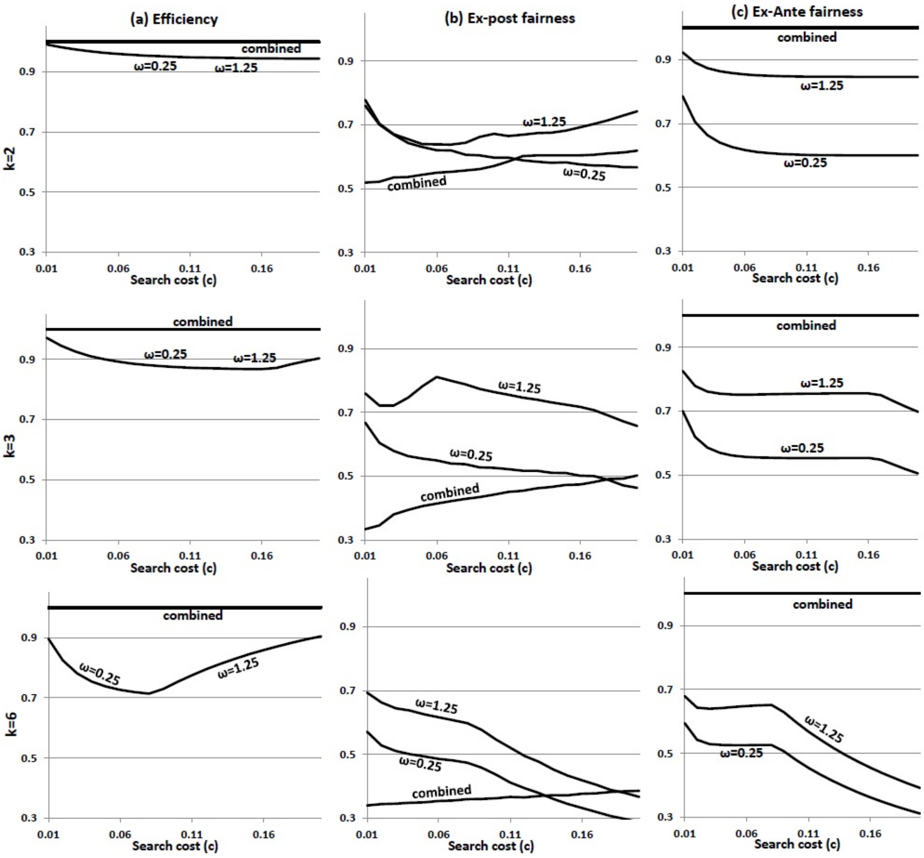

Fig. 7 Trading Efficiency with Fairness: (a) efficiency; (b) ex-post fairness; and (c) ex-ante fairness; using non-efficiency-maximizing policies using $g(x)$ according to (22) with $w = 0.25$ and $w = 1.25$; and the best (i.e., the one associated with the best ex-post fairness) β-based efficiency-maximizing policy. The setting used is: $n = 5$ and $f(x) = 1, \forall 0 \leq x \leq 1$, otherwise $f(x) = 0$ and the number of agents (k) differs between rows.

(b) for $\mathcal{P}_{bonus}(\hat{b})$, $g(y) = \frac{(k-1)y}{k} + \hat{b} \cdot 1(y < r_{bonus})$ where $\cdot 1(y < r_{bonus})$ is the indicator function that returns 1 if $y < r_{bonus}$ and 0 otherwise.
(c) for $\mathcal{P}_{combined}(\hat{s}, \hat{b}, \beta)$, $g(y) = \beta y + \hat{s} + \hat{b} \cdot 1(y < r_{bonus})$.

Figure 7 depicts performance (efficiency, ex-ante and ex-post fairness) achieved by a policy $\mathcal{P}_{function}(g)$ for which:

$$g(x) = \frac{(k-1)x}{k} + \frac{w}{k-1} \cdot \left| E\left(cost^{\mathcal{P}_{split}}(\text{Buyer})\right) - E\left(cost^{\mathcal{P}_{split}}(A_i \neq \text{Buyer})\right)\right| \quad (22)$$

where $\frac{(k-1)x}{k}$ covers the product's cost and the second element approximates the search cost using a coefficient w multiplied by average difference in the expected cost of the buyer and a non-searching agent $A_i \neq$ Buyer. Specifically, the figure depicts the efficiency and fairness measures for two cases differing by the value assigned to the coefficient w $\{w = 0.25, w = 1.25\}$. In addition, we provide in each graph, as a reference, the performance of the fully efficient and fully ex-ante fair policy $\mathcal{P}_{combined}(s, b, \beta)$ 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 and b are computed according to Proposition 6. Each row of graphs corresponds to a different number of agents k. 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 at the expense of the efficiency measure (graph (a)). Therefore the system designer should consider the tradeoff between the two and choose accordingly.

7 Related Work

Historically, efficiency has been the main objective of teamwork and multi-agent systems [53, 7, 9, 20, 48, 42]. Recently, as agents frequently represent human individuals, the importance of *fairness*, as an independent goal, has gained recognition within the MAS literature (see [14, 10], and in particular the review given in [15]). Many works use the "cake cutting" setting as a model for considering fairness in multi-agent systems. The cake cutting model, introduced in [50], 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 (see [6, 37], and [34] for a recent survey). 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 [8, 1]. Accordingly, an important line of research has been in devising algorithms that provide fairness, while optimizing welfare [12, 2, 5, 33]. The cake cutting literature assumes, for the most part, a non-transferable utility setting, wherein agents cannot pay each other.

Other works consider mediated negotiation procedures that support negotiating agents in reaching Pareto efficient and fair agreements, e.g., in bilateral multi-issue negotiation [29] and computational models, which allow agents to find the most desirable solution according to certain definitions of fairness or optimality [17].

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 invest 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.

Costly search of the kind used in this paper is a prevalent theme in MAS [21, 24, 47, 27, 23, 16]. The idea is that agents need to consume some of their resources in order 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 it is related to Markov Decision Processes (MDP) [3, 35] and decentralized Markov Decision Processes [4], 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 [43,40,22,49]. Most of the work in this area has focused on a representative agent, operating on behalf of the team [48,20,30,9,7,45]. However, these assumed that the designated agent is fully cooperative and as such focused on maximizing efficiency, i.e., attempted to extract an optimal search strategy [39,41,44]. 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 [13,18,26,36,19]). In the economic-search literature, however, we are not aware of consideration of fairness.

8 Conclusion

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. Nevertheless, using this policy we can reach a very high precision of ex-post fairness without the need for simulation (using the calculated ex-ante fairness). 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 these types 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 that 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.

In many settings the different team members will be heterogeneous in their search competence. For example, it is possible that one agent is better at searching online stores, while the other is better in searching physical stores. In such cases, the two agents will differ in their search costs and possibly also the price distribution from which they sample. In such settings, both the expected expense and the fairness measures may greatly depend on the choice of which agent is assigned as the buyer. As long as there is just one object to be purchased, this should not be a problem, since the system designer can calculate the efficiency and fairness obtained with each possible buyer, and choose the option that best fits its needs.

When multiple goods need to be purchased, the problem becomes more complex. Indeed, with multiple goods, even the homogenous setting may exhibit counter-intuitive behavior. Consider, for example, ten homogenous agents in need of ten different, but isomorphic, goods (isomorphic in the sense that all goods share the same price distribution). The intuitive fair task allocation is that each agent be assigned one of the goods. This indeed provides ex-ante fairness. However, it may perform poorly on ex-post fairness, as, in actuality, the search costs, as well as the purchase prices for the different goods, may differ substantially. If, additionally, the agents split the purchase costs based on receipts, the resulting strategy becomes highly inefficient. A much better task allocation for this case is that all goods be purchased by one agent, who will be reimbursed using the policies discussed in this paper. Under this task allocation, the aggregation of the risks at a single agent results in lesser variance, and hence higher expected ex-post fairness. The question of fairness in task allocation with multiple goods is thus an interesting research question, which we defer to future research.

Acknowledgments

Preliminary version of this research appears in the proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-14) [38]. The first coauthor was a student at Bar-Ilan University when the research reported in this paper was carried out. This research was partially supported by the ISRAEL SCIENCE FOUNDATION (grant No. 1083/13) and the ISF-NSFC joint research program (grant No. 2240/15).

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Y. Aumann and Y. Dombb. The efficiency of fair division with connected pieces. In *Internet and Network Economics*, pages 26–37, 2010.
2. X. Bei, N. Chen, X. Hua, B. Tao, and E. Yang. Optimal proportional cake cutting with connected pieces. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, pages 1263–1269, 2012.
3. R. Bellman. A Markovian decision process. *Indiana University Mathematics Journal*, 6(4):679–684, 1957.
4. D. Bernstein, D. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27(4):819–840, 2002.
5. S. J Brams, M. Feldman, J. K Lai, J. Morgenstern, and A. D Procaccia. On maxsum fair cake divisions. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, pages 1285–1291, 2012.
6. S.J. Brams and A.D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
7. K. Burdett and D. A. Malueg. The theory of search for several goods. *Journal of Economic Theory*, 24(3):362–376, 1981.
8. I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. The efficiency of fair division. *Theory of Computing Systems*, 50(4):589–610, 2012.

9. J. A. Carlson and R. P. McAfee. Joint search for several goods. *Journal of Economic Theory*, 32(2):337–345, 1984.
10. 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.
11. M. Chhabra, S. Das, and D. Sarne. Expert-mediated search. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 415–422, 2011.
12. Y. J Cohler, J. K Lai, D. C Parkes, and A. D Procaccia. Optimal envy-free cake cutting. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011)*, pages 626–631, 2011.
13. A. Dannenberg, T. Riechmann, B. Sturm, and C. Vogt. Inequity aversion and individual behavior in public good games: An experimental investigation. *ZEW-Centre for European Economic Research Discussion Paper*, 07-034, 2007.
14. S. de Jong and K. Tuyls. Human-inspired computational fairness. *Journal of Autonomous Agents and Multi-Agent Systems*, 22(1):103–126, 2011.
15. S. de Jong, K. Tuyls, and K. Verbeeck. Fairness in multi-agent systems. *The Knowledge Engineering Review*, 23(2):153–180, 2008.
16. A. Elmalech, D. Sarne, and B. J. Grosz. Problem restructuring for better decision making in recurring decision situations. *Autonomous Agents and Multi-Agent Systems*, 29(1):1–39, 2015.
17. U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proceedings of the Second International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2003)*, pages 177–184, 2003.
18. E. Fehr and K. M. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868, 1999.
19. E. Fehr and K. M Schmidt. Theories of fairness and reciprocity-evidence and economic applications. *CEPR Discussion Paper*, 2001.
20. J. Gatti. Multi-commodity consumer search. *Journal of Economic Theory*, 86(2):219–244, 1999.
21. A. Grosfeld-Nir, D. Sarne, and I. Spiegler. Modeling the search for the least costly opportunity. *European Journal of Operational Research*, 197(2):667–674, 2009.
22. B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.
23. C. Hajaj, N. Hazon, D. Sarne, and A. Elmalech. Search more, disclose less. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*, 2013.
24. N. Hazon, Y. Aumann, S. Kraus, and D. Sarne. Physical search problems with probabilistic knowledge. *Artificial Intelligence*, 196:26–52, 2013.
25. M. C. W. Janssen, J. L. Moraga-Gonzalez, and M. R. Wildenbeest. Truly costly sequential search and oligopolistic pricing. *International Journal of Industrial Organization*, 23(5-6):451–466, 2005.
26. D. Kahneman, J. L Knetsch, and R. H Thaler. Fairness and the assumptions of economics. *The Journal of Business*, 59(4):285–300, 1986.
27. J. Kephart and A. Greenwald. Shopbot economics. *Journal of Autonomous Agents and Multi-Agent Systems*, 5(3):255–287, 2002.
28. M. Kohn and S. Shavell. The theory of search. *Journal of Economic Theory*, 9(2):93–123, 1974.
29. M. Li, Q. B. Vo, and R. Kowalczyk. Searching for fair joint gains in agent-based negotiation. In *Proceedings of the 8th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 1049–1056, 2009.
30. E. Manisterski, D. Sarne, and S. Kraus. Enhancing cooperative search with concurrent interactions. *Journal of Artificial Intelligence Research*, 32(1):1–36, 2008.
31. J. McMillan and M. Rothschild. Search. In *Proceedings of Handbook of Game Theory with Economic Applications*, pages 905–927, 1994.
32. P. Morgan and R. Manning. Optimal search. *Econometrica*, 53(4):923–944, 1985.
33. E. Mossel and O. Tamuz. Truthful fair division. In *Proceedings of the 3rd Symposium on Algorithmic Game Theory*, pages 288–299, 2010.
34. A. D Procaccia. Cake cutting: not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.
35. M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience, 1994.
36. M. Rabin. Incorporating fairness into game theory and economics. *The American economic review*, pages 1281–1302, 1993.
37. J. Robertson and W. Webb. *Cake-Cutting Algorithms: Be Fair if You Can*. Ak Peters Series. Taylor & Francis, 1998.
38. I. Rochlin, Y. Aumann, D. Sarne, and L. Golosman. Efficiency and fairness in team search with self-interested agents. In *Proceedings of the 13th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pages 365–372, 2014.

39. I. Rochlin and D. Sarne. Information sharing under costly communication in joint exploration. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*, pages 847–853, 2013.
40. I. Rochlin and D. Sarne. Constraining information sharing to improve cooperative information gathering. In *Proceedings of the 13th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pages 237–244, 2014.
41. I. Rochlin and D. Sarne. Utilizing costly coordination in multi-agent joint exploration. *Multiagent and Grid Systems*, 10(1):23–49, 2014.
42. I. Rochlin, D. Sarne, and M. Laifenfeld. Coordinated exploration with a shared goal in costly environments. In *Proceedings of the ECAI 2012 - 20th European Conference on Artificial Intelligence*, pages 690–695, 2012.
43. I. Rochlin, D. Sarne, and M. Mash. Joint search with self-interested agents and the failure of cooperation enhancers. *Artificial Intelligence*, 214:45–65, 2014.
44. I. Rochlin, D. Sarne, and G. Zussman. Sequential multilateral search for a common goal. In *Proceedings of the 2011 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2011)*, pages 349–356, 2011.
45. I. Rochlin, D. Sarne, and G. Zussman. Sequential multi-agent exploration for a common goal. *Web Intelligence and Agent Systems*, 11(3):221–244, 2013.
46. M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4):689–711, 1974.
47. D. Sarne. Competitive shopbots-mediated markets. *ACM Transactions on Economics and Computation*, 1(3):17:1–17:41, 2013.
48. D. Sarne, E. Manisterski, and S. Kraus. Multi-goal economic search using dynamic search structures. *Journal of Autonomous Agents and Multi-Agent Systems*, 21(2):204–236, 2010.
49. O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, 1998.
50. H. Steinhaus. Sur la division pragmatique. *Econometrica*, 17:315–319, 1949.
51. Z. Tang, M. D. Smith, and A. Montgomery. The impact of shopbot use on prices and price dispersion: Evidence from online book retailing. *International Journal of Industrial Organization*, 28(6):579–590, 2010.
52. R. Waldeck. Search and price competition. *Journal of Economic Behavior and Organization*, 66(2):347–357, 2008.
53. G. Weiss, editor. *Multiagent systems: a modern approach to distributed artificial intelligence*. MIT Press, 1999.
54. M. L. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, 1979.