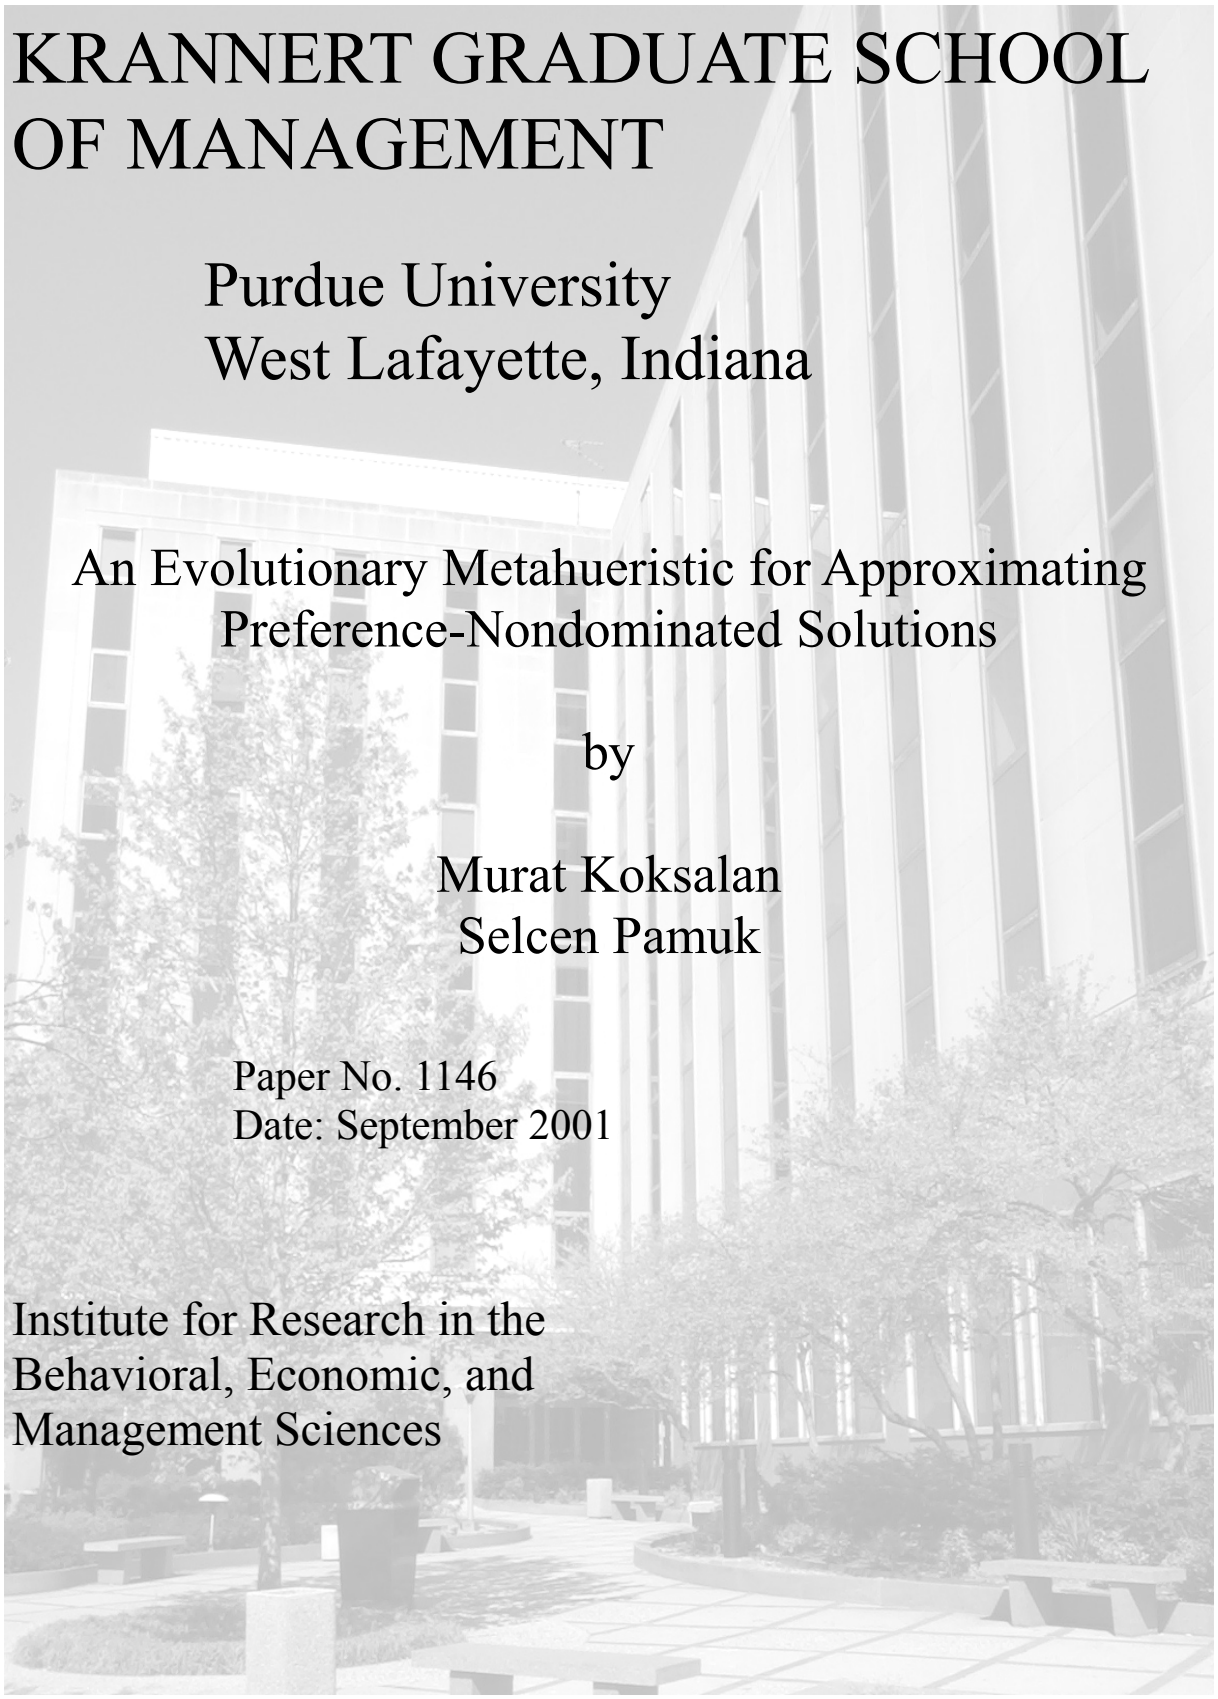

KRANNERT GRADUATE SCHOOL OF MANAGEMENT

Purdue University
West Lafayette, Indiana

An Evolutionary Metahueristic for Approximating Preference-Nondominated Solutions

by

Murat Koksalan
Selcen Pamuk

Paper No. 1146
Date: September 2001

Institute for Research in the Behavioral, Economic, and Management Sciences

An Evolutionary Metaheuristic for Approximating Preference-Nondominated Solutions

Murat Köksalan[1,2] and Selcen Pamuk[1]

[1] *Department of Industrial Engineering, Middle East Technical University*
06531, Ankara, Turkey

[2] *Krannert School of Management, Purdue University*
West Lafayette, IN 47907, USA
E-mail: koksalanm@mgmt.purdue.edu

September 2001

An Evolutionary Metaheuristic for Approximating Preference-Nondominated Solutions

Murat Köksalan[1,2] and Selcen Pamuk[1]

[1]Industrial Engineering Dept., Middle East Technical University, 06531 Ankara, Turkey
[2] Krannert School of Management, Purdue University, West Lafayette, IN 47906, USA, E-mail: koksalanm@mgmt.purdue.edu (correspondence address)

Abstract

We propose an evolutionary metaheuristic for approximating the preference-nondominated solutions of a decision maker in multiobjective combinatorial problems. The method starts out with some partial preference information provided by the decision maker, and utilizes an individualized fitness function to converge towards a representative set of solutions favored by the information at hand. The breadth of the set depends on the precision of the partial information available on the decision maker's preferences. The algorithm simultaneously evolves the population of solutions out towards the efficient frontier, focuses the population on those segments of the efficient frontier that will appeal to the decision maker, and disperses it over these segments to have an adequate representation. Simulation runs carried out on randomly generated instances of the Multiobjective Knapsack Problem and the Multiobjective Spanning Tree Problem have found the algorithm to yield highly satisfactory results.

1. Introduction

Since many real-life problems are multiobjective in nature and have combinatorial components, the last decade has witnessed a growing interest in Multiobjective Combinatorial Optimization (MOCO). The majority of the available MOCO methods concentrate on obtaining an approximation of the efficient frontier of the problem. Even if a good approximation can be

found, the fact remains that the true goal of the decision maker (DM) in a concrete application is a single solution to be implemented, or at most a small subset of solutions for detailed consideration. Concluding the exploration with an approximation of the entire efficient frontier hence necessitates a subsequent search, where the DM will have to first spend some time eliminating irrelevant solutions, and then choose from among a limited selection of relevant ones.

The present study draws on the proven strength of metaheuristics in addressing combinatorial problems and on the work done to date on partial information in the field of Multiple Criteria Decision Making to develop an evolutionary metaheuristic to approximate only the relevant segments of the efficient frontier. The method starts out with some partial information on DM preferences, and evolves a population of solutions towards those regions of the objective space that will appeal to the DM. An individualized fitness function is utilized to allow each solution to represent itself as favorably as possible in light of the known preferences of the DM. By means of this fitness function, the algorithm converges towards, not a single solution, but a set of solutions favored by the information at hand. The breadth of the set depends on the precision of the partial information available. The DM may then either search the final population of the metaheuristic interactively, or draw a sample of diverse solutions that is small enough for his/her detailed consideration. Simulation runs indicate that the algorithm yields highly satisfactory solutions for randomly generated instances of the Multiobjective Knapsack Problem and the Multiobjective Spanning Tree Problem.

The next section overviews the relevant literature and defines the terms that will be used in the rest of the text. Section 3 highlights the significance of the method being proposed and gives the

details of the algorithm. Section 4 summarizes computational findings for test problems, and finally Section 4 presents our conclusions and directions for further research.

2. Background Literature And Terminology

Multiple Objective Decision Making (MCDM) is characterized by the existence of trade-offs between different objectives, which must be evaluated in terms of the preferences of a given DM. In MOCO, this is further complicated by a large and irregular search space causing even the single-objective versions of the underlying problems to be NP-hard. A critical survey of the scope and methods of MCDM is given in Stewart (1992), while surveys of MOCO may be found in Ulungu and Teghem (1994) and Ehrgott and Gandibleux (2000).

In the presence of multiple criteria, the concept of optimality breaks down, to be replaced by that of efficiency. A solution is said to be *efficient* if there exists no other solution performing at least as well in terms of all objectives and better in terms of at least one objective. Together, all efficient solutions to the problem define the *efficient frontier*. Formally, let X be the set of feasible solutions and $f_k(x)$, k = 1, ..., p, be the k^{th} objective function evaluated at solution $x \in X$. We say that $y \in X$ is an efficient solution to the MOCO problem

$$\text{"}\underset{x \in X}{Max}\text{"}\{f_1(x), f_2(x), \ldots, f_p(x)\}$$

if there exists no $x \in X$ such that $f_k(x) \geq f_k(y)$ for all k = 1, ..., p, with at least one inequality strict. If any such x exists, it is said to *dominate* y. The *ideal point* of the problem, $f^* = (f_1^*, f_2^*, \ldots, f_p^*)$, dominates all feasible solutions and is found by separately maximizing the objectives:

$$f_k^* = \underset{x \in X}{Max}\{f_k(x)\}$$

Given a MCDM problem without any information on the preferences of the DM, a particular solution $x \in X$ may be said to be superior to another $y \in X$ if and only if x dominates y. If the DM's preference structure is known and can be represented by a utility function

$$U(x) = u\big(f_1(x),..., f_p(x)\big),$$

then the superiority of solutions will be determined by their appeal to the DM and solution x may be said to be *preferred* to solution y if and only if

$$U(x) \geq U(y).$$

Given the inherent difficulty of expressing and measuring preferences, there will usually be only *partial* or *incomplete information* available on the function U. If we represent all likely utility functions by the set Ω, then we would be able to say that *x is preferred to y under* Ω if and only if

$$U(x) \geq U(y) \quad \forall \quad U \in \Omega.$$

If no such $x \in X$ exists, then solution y is *preference-nondominated under* Ω. Moreover, if a solution y has the highest utility score for at least one likely utility function, i.e.,

$$\exists \quad U \in \Omega \quad \ni \quad U(y) = \underset{x \in X}{Max}\{U(x)\},$$

then y is *potentially optimal.*

An overview of issues related to decision making under partial information on U, and of the early literature on the subject, may be found in Weber (1987). While occasional studies (see, for instance, Hazen, 1986) have considered varying families of utility functions, the widespread tendency has been to assume that U is linear

$$U(x) = \sum_{k=1}^{p} w_k f_k(x)$$

and that partial information pertains to the vector $w = (w_1, \ldots, w_p)$ of *criteria weights*. Takeda and Satoh (2000) have used a Data Envelopment Analysis approach to provide a partial ranking of alternatives under imperfect information on criteria weights for an additive value function. Most other work in the field has aimed to identify preference-nondominated or potentially optimal solutions (see, for instance, Athanassopoulos and Podinovski, 1997).

If the DM is able to formulate linear inequalities constraining the possible values of w, then these inequalities, in conjunction with a normalizing constraint of the form

$$\sum_{k=1}^{p} w_k = 1$$

form a polytope W of all possible weight vectors. Under these circumstances, solution x is at least as preferred as solution y under W if and only if

$$\sum_{k=1}^{p} w_k \left(f_k(x) - f_k(y)\right) \geq 0 \quad \forall \quad w \in W .$$

Special cases where W's characteristics simplify checking whether the above is true have been addressed by Kirkwood and Sarin (1985) and Carrizosa et al (1995).

Malakooti et al. (1994) have distinguished between the two cases where (i) there exists an underlying vector $w^{act} \in W$ defining U which cannot be more precisely determined, and where (ii) several different vectors in W may be valid for different points in X, in effect rendering the utility function U nonlinear. The significance of the distinction lies in the potential optimality of preference-nondominated solutions that are convex dominated. An alternative $y \in X$ is said to be *convex dominated* by the alternatives in the set $X \backslash \{y\}$ if at least one convex combination of the alternatives in this set dominates y, i.e.,

$$\exists \ \mu_i \ \ni \ \sum_{x^i \in X \backslash \{y\}} \mu_i f_k(x^i) \geq f_k(y) \quad \forall \quad k = 1,..,p; \qquad \sum_{x^i \in X \backslash \{y\}} \mu_i = 1; \qquad \mu_i \geq 0 \quad \forall \quad x^i \in X \backslash \{y\}$$

In case (i), it is impossible for convex dominated alternatives to be potentially optimal, but in case (ii), these solutions also may be potentially optimal.

Extensions of the work done on MCDM under partial preference information to MOCO are virtually nonexistent. The majority of authors studying MOCO problems have aimed to characterize or approximate the full efficient frontier, although the cardinality of this set may be exponential in problem size. Occasional interactive metaheuristics have been designed that progressively elicit information from an accessible DM and converge towards the unique solution that represents the DM's most favored compromise of objectives (Teghem, Tuyttens, and Ulungu, 2000; Pamuk and Köksalan, 2001). The partial information case where the DM is not accessible beyond providing some limited preference information *a priori*, on the other hand, has to our knowledge not been addressed.

The success of multiobjective extensions of metaheuristics such as Tabu Search (for instance, Gandibleux, Mezdaoui, and Freville, 1996), Simulated Annealing (for instance, Czyzak and Jaszkiewicz, 1998), and Evolutionary Algorithms (for instance, Zitzler and Thiele, 1999) in approximating the efficient frontiers of MOCO problems makes it likely that these methods may further be extended to work in the partial information case.

3. Evolutionary Metaheuristic for Approximating Preference-Nondominated Solutions

The significance of the Evolutionary Metaheuristic for Approximating Preference-Nondominated Solutions (EMAPS) proposed here lies in the fact that all available preference information is used to guide and restrict the search effort. Instead of spending computational resources securing the entire efficient frontier, some parts of which are of no interest to the DM, the portions that are of interest are searched in greater depth. The subsequent search for a single solution is hence

carried out in a narrower and more relevant search space. The more precise the partial information on the DM's preferences, the narrower will be the space in question. The algorithm is hence a flexible tool, and may be used for all cases between the two extremes of perfect information and no information. If the utility function of the DM is known precisely, then the algorithm will approximate the optimum for that function; whereas in the absence of any preference information, the full efficient frontier will be approximated.

Figure 1 below compares the evolution directions of EMAPS with those of other methods in literature. Traditionally, the goals of the multiobjective evolutionary metaheuristics are twofold: (1) to push the population of solutions out towards the efficient frontier, and (2) to disperse it over the frontier to have an adequate representation. Dominance-based fitness functions are the most common tool for guiding the population in direction (1), while direction (2) is usually achieved by some form of fitness sharing or niching. In EMAPS, a third goal has been added to the other two, as the algorithm simultaneously strives (3) to focus the population on those segments of the efficient frontier that will appeal to the DM, given our partial information on his/her preferences.

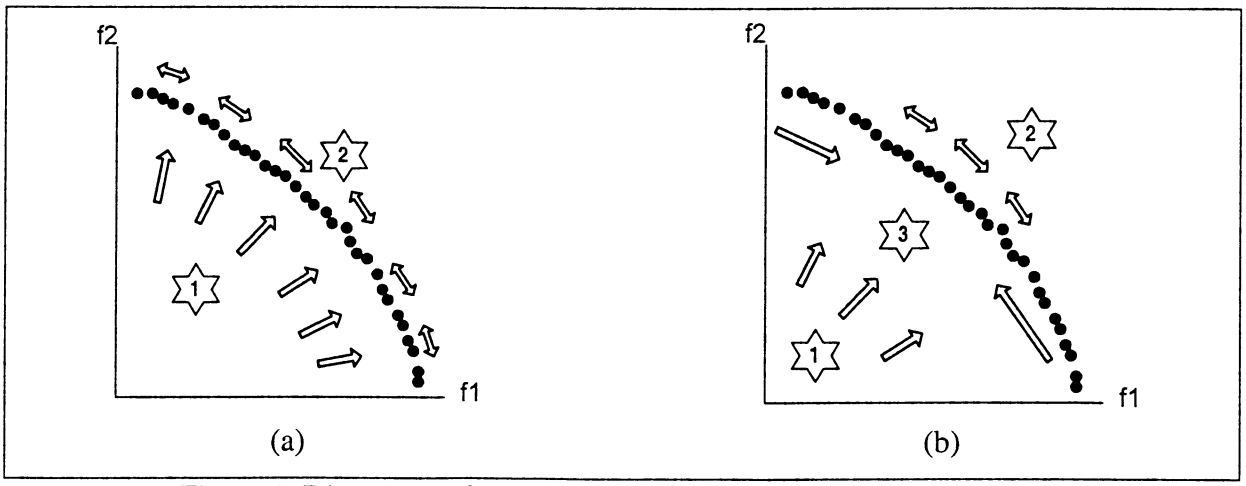

Figure 1. Directions of evolution for (a) traditional methods (b) EMAPS

EMAPS evolves a population of solutions with the goal of generating a good approximation of the DM's preference-nondominated solution set (X^{PND}) under a constrained weight set W. Although it is assumed that W is defined by linear inequalities and consistent with a linear utility function, the algorithm does not automatically eliminate convex dominated solutions. A major factor in the design of the algorithm is the use of strategies that encourage the population's diversity and hence a good coverage of the entire set X^{PND}. For this purpose, duplicate solutions are not permitted to be simultaneously present in the population. The fitness score of each new offspring is calculated by solving an LP model. This allows each member of the population to be evaluated in terms of the weight vector $w \in W$ that puts it at the greatest advantage, and assigns higher scores to those solutions that approximate the portions of X^{PND} that are underrepresented. Finally, new solutions entering the population tend to replace inferior solutions to which they are similar.

3.1 The Restricted Weight Space

For the purposes of this study, we will assume that the DM provides us with some information that can be transformed into linear inequalities on criteria weights, which we use to constrain the weight space W. The information may include, for instance, a partial ranking of objectives, pairwise comparisons of solutions, and upper and lower bounds on objective weights. If the criteria have been scaled to have approximately equal ranges of values, then statements on the relative importance of criteria may be directly translated into inequalities on criteria weights. Statements of the form:

- $f_k \succ f_m$ (the DM finds objective k to be at least as important as objective m)
- $x^i \succ x^j$ (the DM prefers solution x^i over solution x^j)

- $w_k^{LB} \le w_k \le w_k^{UB}$ for a subset of objectives $K \subseteq \{1, \ldots, p\}$

may thus be utilized to form the set:

$$W = \left\{ w = (w_1, w_2, \ldots, w_p) \;\middle|\; \begin{array}{c} w_k \ge w_m \quad \forall \quad f_k \succ f_m \\ \sum_{k=1}^{p} w_k \left(f_k(x^i) - f_k(x^j) \right) \ge 0 \quad \forall \quad x^i \succ x^j \\ w_k^{LB} \le w_k \le w_k^{UB} \quad \forall \quad k \in K \\ \sum_{k=1}^{p} w_k = 1 \\ w_k \ge 0 \quad \forall \quad k = 1, \ldots, p \end{array} \right\}$$

We assume that W is a nonempty set, i.e., that the information provided by the DM is consistent, both internally and with our assumptions on the form of the utility function. The resolution of inconsistencies during the specification of W remains a problem of some interest, but is not addressed here, since various resolution methods may conceivably be devised without much difficulty.

3.2 Forward Filtering

At several points in the algorithm, a set of p-dimensional (solution or weight) vectors is filtered to obtain a smaller diverse set by an operation known as *forward filtering*. The idea is to choose those vectors that are approximately evenly spaced and collectively cover all parts of the space covered by the entire set. This is achieved by constructing a subset of elements that are as dispersed as possible with respect to a given metric, and is explained in detail in Steuer, 1986 (pp.311-321). In EMAPS, Euclidean distances have been selected as the metric to be maximized in constructing the forward-filtered subset.

3.3 Initialization and Weight Space Sampling

The initial population of EMAPS consists of some solutions generated randomly, and some relatively higher-quality solutions used to seed the population. The seed solutions are produced by repeated applications of a base heuristic known to work well on the single-objective version of the underlying MOCO problem; where the single-objective problems to be solved are formed by linearly aggregating the objective functions using different weight vectors obtained through an even sampling of the restricted weight space W.

In the absence of any special structure for W, it is possible to generate a large set R of weights by incrementing each w_k uniformly from 0 to 1 by some step size 1/r. The set

$$R = \left\{ w = (w_1, w_2, ..., w_p) \;\middle|\; \begin{array}{l} w_k \in \left\{0, \frac{1}{r}, \frac{2}{r}, ..., \frac{r-1}{r}, 1\right\} \quad \forall \quad k = 1, ..., p \\ \sum_{k=1}^{p} w_k = 1 \end{array} \right\}$$

formed in this way has cardinality

$$|R| = \binom{r+p-1}{p-1}.$$

The weight vectors belonging to the set W∩R may then be used for the linear aggregation of objectives. If the set W∩R has fewer elements than the desired number of seed solutions, the process may be repeated for a higher value of r. If the set W∩R has too many elements, it may be forward-filtered to the desired cardinality. It should however be noted that not all aggregation weights need lead to distinct solutions, so that the number of seed solutions that are actually obtained may be smaller than the number of weight vectors in the filtered set.

3.4 The Fitness Function

The likelihood that a solution will survive and reproduce during evolution is determined by its fitness score, which, in EMAPS, should reflect the solution's contribution to approximating X^{PND}. One way of evaluating this contribution involves allowing each solution to select a favorable search direction in W, and measuring how superior it is to the other solutions in the population with respect to performances along that direction. Such a fitness function evolves the population in the three directions illustrated in Figure 1: (1) A given solution immediately has higher fitness than any solutions it dominates, so evolution is towards the efficient frontier. (2) Solutions approximating relatively underrepresented portions of X^{PND} have higher fitness than those that are clustered together, since the former outperform a higher proportion of the population under their preferred search direction, while the latter have to vie with each other for fitness. This implicit sharing effect spreads the population over X^{PND}. (3) Solutions approximating X^{PND} have higher fitness than solutions approximating other portions of the efficient frontier that are remote from the search directions indicated by W, so evolution 'focuses' on the available preference information.

3.4.1 Relative strength of solutions

Given the existing population X' at any point in the algorithm, suppose that we would like to compare a new solution x^{new} with the members of X'. Let us denote by $\phi(w, x)$ the linear combination under weight vector w of the criteria values for solution $x \in X$:

$$\phi(w,x) = \sum_{k=1}^{p} w_k f_k(x)$$

If we can somehow compute a favorable weight vector for x^{new}, $w^{new} \in W$, then a measure of x^{new}'s strength relative to some $x^i \in X'$ would be

$$\varepsilon^i = \phi\left(w^{new}, x^{new}\right) - \phi\left(w^{new}, x^i\right),$$

indicating the surplus utility the DM might derive from choosing x^{new} over x^i if s/he indeed has a linear utility function defined by the weights w^{new}. The strength of x^{new} relative to the entire population X' may then be measured by either the average of these differences:

$$\bar{\varepsilon} = \frac{\sum_{x^i \in X'} \varepsilon^i}{|X'|}$$

or the difference between x^{new} and its closest contender along the search direction w^{new}:

$$\varepsilon = \underset{x^i \in X'}{Min}\left\{\varepsilon^i\right\}$$

A negative ε value indicates that x^{new} is outperformed by some $x^i \in X'$ even at its favorable weights, so that it cannot be potentially optimal for linear utility functions. However, it is possible for a convex-dominated solution with $\varepsilon < 0$ to be potentially optimal if the underlying utility function is not linear. By itself, the worst-case measure ε may be a poor indicator of x^{new}'s fitness, as it is possible for an x^{new} with a few close contenders to fall in a fairly underrepresented portion of X^{PND}, and hence outperform many other solutions in X' under its favorable weights. Taken together, the two terms ensure that the fitness function forces the population to maximize not only the weighted criterion values $w_k f_k(x)$, but also the spread of solutions across the relevant portion of the efficient frontier. Thus, in EMAPS, the fitness score is a convex combination of ε and $\bar{\varepsilon}$:

$$fitness\left(x^{new}\right) = \alpha\bar{\varepsilon} + (1-\alpha)\varepsilon$$

for some $\alpha \in [0,1]$. Extensive test runs suggest that α levels of 0.00, 0.25, and 0.50 tend to work well. If computational resources can be spared, a conservative approach would be to make several runs of EMAPS under different values of α and combine the results.

3.4.2 Favorable weight computation

A straightforward application of the above idea is to solve the following LP to find the *favorable weights* w^{new} of x^{new} as the weights that maximize x^{new}'s fitness:

$$
\begin{aligned}
Max \quad & \frac{\alpha}{|X'|} \sum_{x^i \in X'} \varepsilon^i + (1-\alpha)\varepsilon \\
st \quad & \sum_{k=1}^{p} w_k^{new} \left(f_k(x^i) - f_k(x^{new}) \right) + \varepsilon^i = 0 \quad \forall \quad x^i \in X' \\
& \varepsilon - \varepsilon^i \leq 0 \quad \forall \quad x^i \in X' \\
& w^{new} \in W, \\
& \varepsilon \text{ and the } \varepsilon^i \text{ unrestricted in sign}
\end{aligned}
$$

However, since the number of constraints in the above LP is proportional to the cardinality of the population, and since it has to be solved frequently, the calculation of favorable weights and fitnesses may become too time-consuming for large population sizes. As an alternative, it is possible to store a smaller set X" of dispersed high-quality solutions and substitute X" for X' in the LP to be solved for the computation of w^{new}. In our applications of EMAPS, we used as X" a forward-filtered subset of the seed solutions in the initial population. The smaller LP for each x^{new} can then be solved in less time, resulting in substantial time savings over the course of evolution. The objective function of the smaller LP, however, does not directly yield x^{new}'s fitness. Once the favorable weights w^{new} are found, the score fitness(x^{new}) must instead be calculated separately using the entire population X', as explained in the previous section.

3.4.3 Fitness updates

Since the fitness of solution x^{new} is calculated with respect to the population at the time of x^{new}'s entry, fitness values may become outdated as the population evolves. Hence a solution may retain the high fitness score assigned at the time of its entry, and yet be outperformed by many others that entered the population after it. Conversely, if a certain portion of X^{PND} becomes relatively less represented over time, the fitnesses of the solutions falling in that portion may become too low to reflect their increased contribution to approximating X^{PND}. To prevent such occurrences, the fitness values of all solutions in the population are updated upon every entry into, or exit from, the population. This process is not very time consuming since, with the use of X", the favorable weight vectors of solutions do not depend on the members of X' at any given time.

3.5 Evolution

Evolution involves iteratively choosing members of the population to be recombined to form new solutions, which may then replace relatively inferior members of the existing population. While the genetic representation of solutions and the (crossover, mutation, repair, etc) operators being employed during evolution are problem-specific, some generic features of the algorithm are detailed below.

3.5.1 The selection of parents and creation of offspring

The population maintained by EMAPS is always sorted in order of decreasing fitness scores. The first parent for a crossover is then selected in reference to ranks in the sorted population. The probability that the member with rank r will be chosen to become the first parent in the next crossover is given by (see Reeves, 1993):

$$P\left(x^{[r]}\right)=\frac{2(n-r+1)}{n(n+1)}.$$

After the second parent is selected on a purely random basis, crossover, repair, and mutation operators appropriate for the problem at hand are utilized to produce one or more offspring.

3.5.2 Sampling of weights to guide evolutionary operators

It is often possible to design evolutionary (recombination, repair, etc.) operators that exploit a known base heuristic for the single-objective version of the MOCO problem being addressed. If such operators are to be employed by the algorithm, a rule must be specified for the selection of the operator-guidance weight vector w^{og} used to aggregate the objectives and create the single-objective problem to be solved. In EMAPS, the existence of favorable weights proves useful in the choice of w^{og}. For instance, in a recombination operator such as the one described below for multiobjective spanning tree problems, w^{og} may be set equal to the favorable weights of either of the parents, or to some combination of the two. In some cases, time may be gained by storing information on single objective problems obtained under a predetermined set of weight vectors, W^{pre} forward filtered from a larger set $W \cap R$. The member of W^{pre} to be used as w^{og} in a given operation may then be found by minimizing some distance metric from favorable weights.

3.5.3 Population entry and replacement rules to encourage diversity

If a new offspring has a fitness score that is worse than that of the lowest-ranking solution in the population, it is immediately discarded and referred to as a "stillborn" baby. Solutions that are duplicates (in the objective space) of the existing members of the population are also denied entry, although they are not counted among the stillborn. If the new solution is distinct from the existing population members and not stillborn, then it is inserted into the population at the place indicated by its fitness score.

As the new solution x^{new} enters the population, it is compared to each of the existing members $x^i \in X'$ with respect to their own favorable weights. If any x^i performs worse under its own favorable weights w^i than x^{new}, i.e. we have

$$\zeta^i = \phi(w^i, x^i) - \phi(w^i, x^{new}) < 0,$$

then $x^i \in X'$ is a candidate for removal from the population. (While ε^i measures the surplus utility the DM would derive from choosing x^{new} over x^i if s/he had a linear utility function with weights w^{new}, ζ^i measures the surplus utility s/he would derive from choosing x^i over x^{new} if s/he had a linear utility function with weights w^i.)

We set

$$\zeta^t = \underset{\zeta^i<0}{Min}\{\zeta^i\}$$

and remove the solution $x^t \in X'$ with the most negative ζ^i from the population. The fact that x^{new} outperforms x^t under the weights w^t implies that the search directions favoring the two solutions are not very distant, and that x^{new} is more promising than x^t in its ability to approximate X^{PND}. In this sense, the replacement scheme described above tends to replace population members by similar and better-performing ones, and is therefore expected to encourage diversity in the population.

The population cardinality in EMAPS is upward flexible with a predetermined upper limit. So long as the population size is below this limit, we employ only the replacement policy described above, and increase cardinality by one if no solution with negative ζ^i exists. After the upper limit is reached, a second replacement policy is put into effect: in the absence of solutions with

negative ζ^i, the entering solution replaces the member of the population with the lowest fitness score to keep the population size stable.

3.6 Termination

EMAPS may be terminated when any of the following conditions holds:

(i) a predetermined number of new offspring are consecutively stillborn, indicating that the population cannot be improved much further

(ii) the population has converged to a predetermined extent, as measured by the percentage of genes that have the same value in a predetermined majority of the population

(iii) a predetermined number of crossovers have been realized.

Following termination, the population is purged of dominated solutions and forward filtering may be applied to present the DM with a predetermined number of diverse solutions. Alternatively, the DM may choose to have access to the entire final population, e.g., to perform an interactive search on its members.

3.7 Customization for a Specific Application

A flowchart of EMAPS may be found in Figure 2 below. To customize EMAPS to fit a given MOCO problem, it is necessary to specify

(i) a genetic representation of solutions

(ii) crossover and mutation operators and probabilities

(iii) a repair operator or penalty function to be employed if the genetic operators may yield infeasible solutions

(iv) a base heuristic for the single-objective version of the problem

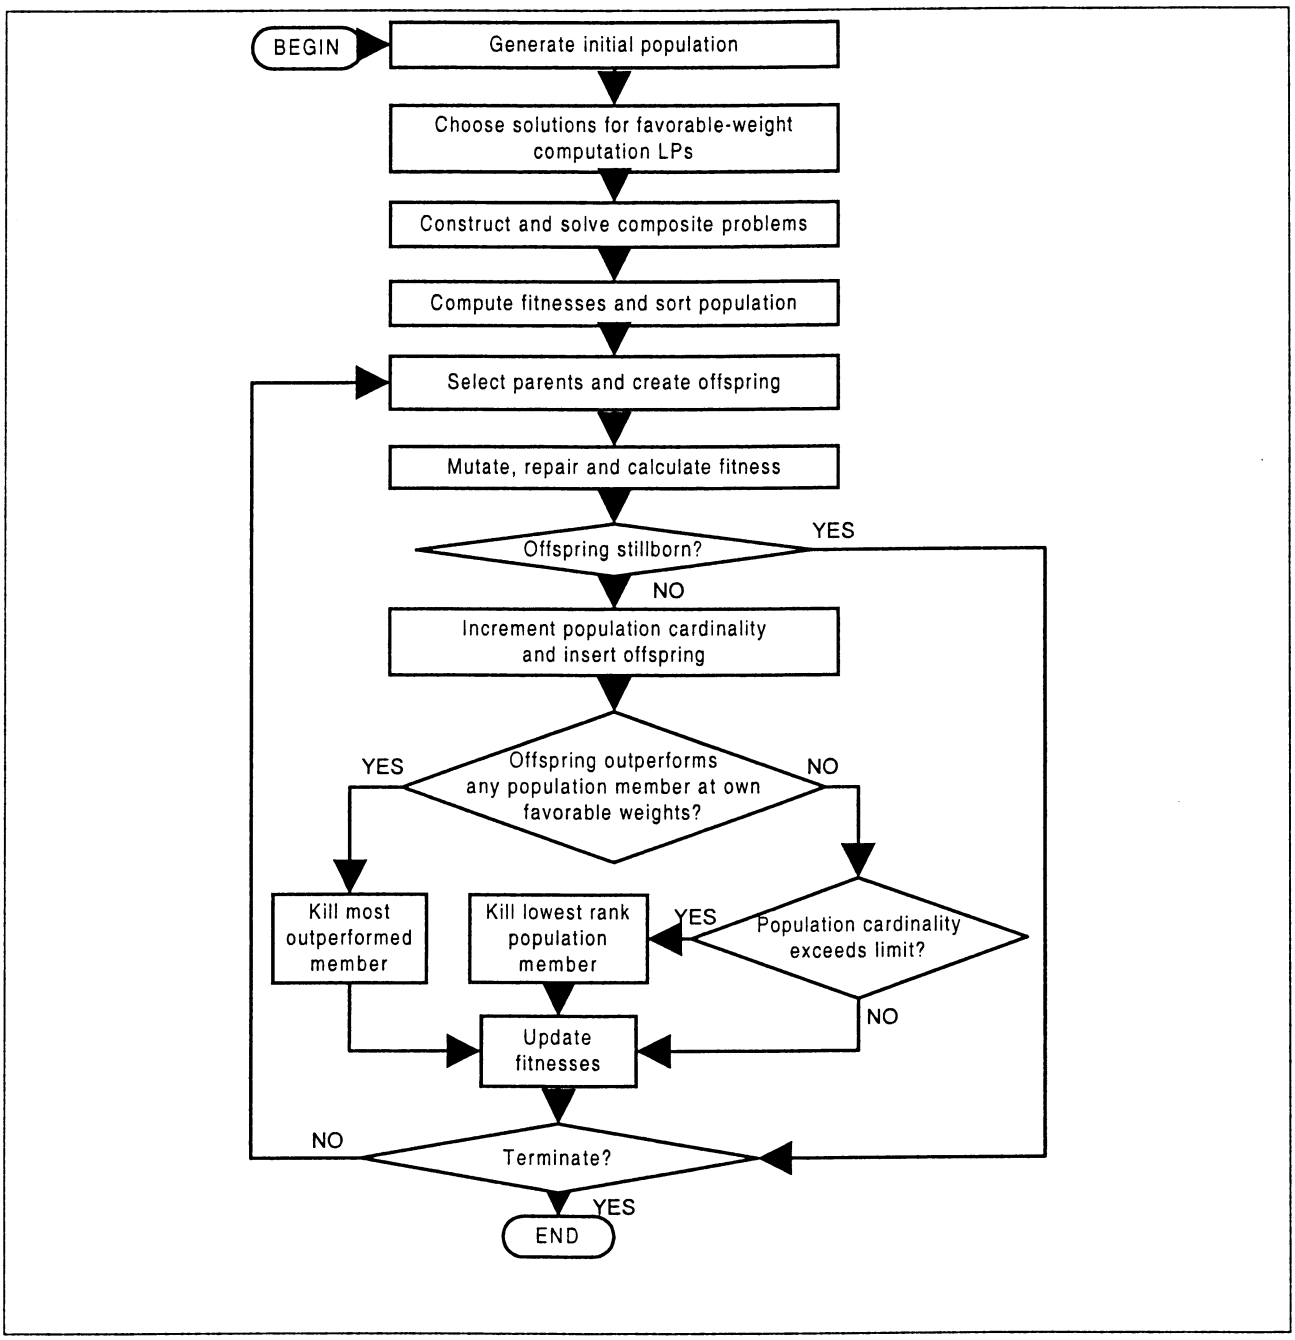

Figure 2. Flowchart for EMAPS

(v) initial and maximum population cardinality values

(vi) number of desired seeds in the initial population

(vii) number of desired weights in W^{pre}

(viii) number of desired solutions to be used in favorable weight LPs

(ix) a value for α, determining the proportion of ε to $\bar{\varepsilon}$ in the fitness definition

(x) termination conditions

(xi) a rule for the selection of a subset of solutions to be presented to the DM or a method for the DM to search the ending population.

4. Computational Findings

To date, the generic metaheuristic described above has been implemented on two well-known MOCO problems, namely the Multiobjective Knapsack Problem (MOKP) and the Multiobjective Spanning Tree Problem (MOST). In all simulation runs, the maximum population cardinality was set to 500 and the maximum number of crossovers to 5000. For larger problems, with broad search spaces and wide W sets, we found these values to be too low, but restrictions on computer resources prevented us from increasing them any further.

Initial population sizes and characteristics were defined with reference to p, since the search space grows exponentially with the number of objectives. Table 1 below indicates the step sizes 1/r used to define the sets R. The targeted number of weight vectors forward filtered out of $W \cap R$ and used to seed the population equaled two thirds of the initial population cardinality.

Table 1. Initial population characteristics for EMAPS

p	2	3	4
Initial population cardinality	60	90	150
Targeted number of seed weights	40	60	100
Step size 1/r	1/800	1/145	1/76
Cardinality of set R	801	10731	79079

The seed solutions were further forward filtered to obtain twenty dispersed high-quality solutions to be used in favorable weight LPs. In addition, W^{pre} weight sets of cardinality 10 were forward filtered out of $W \cap R$ and the corresponding solutions used to guide the repair operator in MOKP and the crossover operator in MOST. To counteract the high intensification tendencies inherent in these operators, the mutation probability was set to 0.90. Crossover probability, on the other hand, was specified as 1.00. The algorithm was terminated when either the maximum number of crossovers was realized, or the population converged (meaning that 95% or more of the genes had the same value in 95% or more of the population), or 50 consecutive offspring were stillborn.

The alpha (α) parameter in the fitness function, measuring the weight of the average versus the minimum utility surplus, was varied to equal α = 0.00, 0.25, 0.50, 0.75 and 1.00 and the performances of each of these versions as well as the union of their outputs were evaluated. The union operation involved first combining the final populations and then deleting duplicated or dominated solutions. As a final point, the performances of forward filtered subsets of the union sets were also investigated.

To evaluate EMAPS's performance, it is necessary to recall that the algorithm aspires to provide the DM with a good approximation of his/her preference-nondominated solution set (X^{PND})

under the constrained weight set W. If there is actually a single element of W defining the DM's preferences, which cannot be more precisely determined *a priori*, then the performance measure must be in reference to this single underlying utility maximization problem. If, on the other hand, all elements of W are equally likely to define the DM's preferences, or perhaps different weight vectors are valid for different regions of the efficient frontier, then the performance measure must take the entire W into account.

Our simulation runs on MOKP and MOST involved actual criteria weights w_k^{act} randomly generated from a discrete uniform distribution over [0, 1] and then normalized to sum to one. The utility maximization problem

$$\underset{x \in X}{Max} \; U(x)$$

was then solved for the utility functions

- Linear $\quad U(x) = -\sum_{k=1}^{p} w^{act}{}_k \left(f_k^* - f_k(x)\right)$
- Tchebycheff $\quad U(x) = -\underset{k=1,\ldots,p}{Max}\left\{w^{act}{}_k \left(f_k^* - f_k(x)\right)\right\}$

where f* is the ideal point for the problem. Linear and Tchebycheff utility functions are commonly used for simulations in literature (see, for example, Köksalan and Sagala, 1995) because they correspond to the opposing DM tendencies of aggregation versus balance of objectives. Performance may hence be evaluated by the scaled deviation percentages

$$\delta = 100 \frac{U\left(x^{BENCH}\right) - U(x^{BEST})}{U\left(x^{BENCH}\right) - U(x^{BAD})},$$

where x^{BENCH} is a good benchmark solution to the utility maximization problem, x^{BAD} is a poor solution, and x^{BEST} is the solution with highest U(x) value from among all the solutions in the final population of EMAPS or a subset thereof presented to the DM.

In addition, W sets of varying sizes were constructed around the normalized weights w^{act} as:

$$W = \left\{ w = (w_1, w_2, ..., w_p) \;\middle|\; \begin{array}{c} \sum_{k=1}^{p} w_k = 1 \\ w_k^{act} - \omega \le w_k \le w_{act}^k + \omega \quad \forall \quad k = 1, ..., p \end{array} \right\}$$

for the values ω = 0.05, 0.15, and 0.25. The W sets corresponding to these three ω levels were referred to as the cases of W1, W2, and W3, respectively. The bound constraint on w_k was replaced by $0 \le w_k \le 2\omega$ whenever we had w_k^{act} - ω < 0 and by $1 - 2\omega \le w_k \le 1$ whenever we had $w_k^{act} + \omega > 1$.

The weight vectors falling in the W∩R sets were used (without filtering) with linear utility functions in the generation of approximate preference-nondominated solution sets $\hat{X}^{PND}$. These sets are approximate in the sense that they do not contain any convex dominated solutions and may not contain all supported efficient solutions in the range since not all possible w∈W are used, although the sampling of W is quite dense.

Scaled deviation percentages $\delta|\hat{x}$ were calculated for each $\hat{x} \in \hat{X}^{PND}$ and the average

$$\delta_{avg} = \frac{\sum_{\hat{x} \in \hat{X}^{PND}} \delta|\hat{x}}{\left|\hat{X}^{PND}\right|}$$

was used to measure how successfully EMAPS approximates X^{PND} on the average. Furthermore, the $\delta|\hat{x}$ value for the solution $\hat{x}$ that is approximated with the least success,

$$\delta_{max} = \underset{\hat{x} \in \hat{X}^{PND}}{Max} \{\delta|\hat{x}\},$$

measured representative ability by identifying any gaps that may exist in the population.

4.1 PDEM Implementations on MOKP

Given the set {1, ..., J} of items that may be placed into a knapsack of capacity b, the Multiobjective Knapsack Problem (MOKP) is formulated as:

$$\text{"}Max\text{"} \left\{ \sum_{j=1}^{J} c_j^k x_j, \quad k = 1,.., p \right\}$$
$$st. \sum_{j=1}^{J} a_j x_j \le b$$
$$x_j \in \{0,1\} \quad \forall j = 1,..., J$$

where c^k_j is the contribution of the j^{th} item to the k^{th} objective, a_j is the volume of the j^{th} item, and x_j is a binary variable that takes on the value 1 if the j^{th} item is included in the solution and 0 otherwise. For each of three problem dimensions given by p = 2, 3, and 4, ten different instances of MOKP were generated by varying the random number seeds used in the generation process. For all instances, J was fixed at 200, c^k_j and a_j values were randomly generated from a discrete uniform distribution over [60, 100], and knapsack capacity b was set to half of the total item volumes in an effort to make the problem harder to solve (see Martello and Toth, 1990).

The base single-objective heuristic used to seed the initial population and to perform local search in later iterations is the greedy algorithm, which fills the knapsack from scratch in order of decreasing contribution per unit capacity usage ratios c_j/a_j. The greedy algorithm was also used under actual weights w^{act} in reverse, by filling the knapsack in order of increasing ratios, to obtain the x^{BAD} solutions. The x^{BENCH} solutions in reference to the single utility functions defined by w^{act} were found by an optimization package. However, since the combined cardinality of all the R sets equaled 906,110, solving all problems induced by the weight vectors in the $W \cap R$ sets proved computationally too excessive. The $\hat{X}^{PND}$ sets for MOKP were, therefore, formed by

solving linear relaxations of the utility maximization problems, and it was found that even the δ_{avg} and δ_{max} values computed in reference to these 'upper bound' solutions were quite excellent.

Straightforward genetic representations and crossover/mutation operators were adapted from single-objective knapsack literature (Beasley and Chu, 1998) as follows: Each solution is represented by a binary chromosome, each gene of which holds the value of one decision variable x_j. The genetic operators chosen are *uniform crossover* and *single bit flip mutation.* Hence, given the parent solutions x^{P1} and x^{P2}, a random binary *crossover mask* of length J is generated, and the offspring solutions x^{C1} and x^{C2} are created by interpreting 0's in the mask as marking the genes x^{C1} will inherit from x^{P1}, and x^{C2} from x^{P2}. Similarly, a 1 in the crossover mask marks a gene that will be inherited from x^{P2} by x^{C1} and from x^{P1} by x^{C2}. Any offspring solution that is not stillborn is then mutated with a predetermined probability. Mutation involves the random selection of one of J genes, which is then flipped from 0 to 1 or 1 to 0, corresponding to the insertion or removal of an item, respectively.

Since the genetic operators outlined above do not guarantee feasibility, a greedy empty-and-fill repair operation is defined as follows: At initialization, the vectors $w^{og} \in W^{pre}$ are used in turn to define a composite objective function coefficients

$$c_j = \sum_{k=1}^{p} w_k^{og} c_j^k$$

and the items sorted with respect to their contribution per unit capacity usage ratios c_j/a_j. An additional binary variable in uniform crossover determines which offspring will be assigned the favorable weights of which parent. Then, when an offspring assigned the weight vector w^{asg} is to be repaired, the sorted item list corresponding to the $w^{og} \in W^{pre}$ with the shortest Euclidean distance from w^{asg} is used. If a solution violates the knapsack capacity, items are taken out of the

knapsack in the list order until feasibility is restored. If a solution (either upon generation or after an empty operation) does not utilize the full knapsack capacity it could command, then additional items are put into the knapsack in the list order.

4.2 PDEM Implementations on MOST

The Multiobjective Spanning Tree (MOST) problem on a graph G = (N, E) with edge weights $c_k(e)$, k=1,...,p, associated with each edge e∈E, is formulated as

$$\text{"}\underset{T\in\tau(G)}{Max}\text{"}\left\{-\sum_{e\in T} c_k(e),\quad k=1,..,p\right\},$$

where τ(G) denotes the set of all spanning trees of G. For the purposes of this study, twenty instances of the MOST problem were generated for each value of p = 2, 3, 4. The first ten problems in each dimension (MOST20) were over complete graphs on 20 nodes; the other ten problems (MOST50) over complete graphs on 50 nodes. In both cases, $c_k(e)$ values were randomly generated from a discrete uniform distribution of range [60, 100].

A greedy algorithm adding low-cost edges to a growing tree so long as cycles are not created was used, not only to seed the initial population, but also to find x^{OPT} solutions to the problems with linear utility functions. Moreover, the greedy algorithm was used in reverse to obtain x^{BAD} solutions. The x^{OPT} solutions for the problems with Tchebycheff utility functions were sought by Hamacher and Ruhe's (1994) procedure evaluating the 9000 best minimum spanning trees. Even for those cases when this procedure did not result in an optimal solution, it yielded x^{GOOD} solutions with excellent bounds on relative accuracy. Table 2 below gives the number of problems (out of ten) where x^{OPT} was found, and for those problems where it could not, indicates the maximum distance of x^{GOOD} from optimality, as a percentage of the bound on the objective

Table 2. Benchmarks for MOST problems with Tchebycheff utility function

Problem	p	Cases where x^{OPT} was found	Max. distance of x^{GOOD} from optimality
MOST20	2	10	-
	3	7	0.09%
	4	0	0.31%
MOST50	2	8	0.05%
	3	7	0.06%
	4	0	0.30%

value. Finally, the $\hat{X}^{PND}$ sets approximating the DM's preference nondominated solutions were found by using the greedy algorithm to solve linear utility maximization problems to optimality. No genetic representations were utilized and the *cut-and-combine mutation* and *subgraph-MST crossover* operators we designed dealt directly with the spanning trees. Subgraph-MST crossover is based on the greedy algorithm and finds the single-objective Minimum Spanning Tree of the subgraph obtained by taking only the edges of the parent trees into consideration. At initialization, sorted arc lists are formed and stored for ten different composite arc cost values corresponding to $w^{og} \in W^{pre}$. Then, when parent solutions x^{P1} and x^{P2} are to be recombined, the subgraph-MST crossover operator is applied twice to the subgraph formed by juxtaposing the arcs of x^{P1} and x^{P2}, yielding two offspring solutions x^{C1} and x^{C2}. The composite arc costs used to create x^{C1} are those found for the member of W^{pre} with shortest Euclidean distance from x^{P1}'s preferred weights w^{P1}, while those used to create x^{C2} correspond to the member of W^{pre} with shortest Euclidean distance from w^{P2}. After the offspring have been formed, the mutation operator deletes a random arc in the spanning tree and randomly inserts another arc that would combine the two resulting trees into another spanning tree. Hence, the genetic operators being used always yield feasible solutions and no repair operator or penalty function is required. The

mutation operator is rather disruptive, as it is hoped that this will counterbalance the strong intensification tendency inherent in the crossover operator.

4.3 Simulation results

Tables 3 and 4 below present the average of the δ_{avg} values and the maximum of the δ_{max} values observed over the ten problem instances of each problem type. In each row, the minimum value for the EMAPS versions corresponding to the α levels of 0.00, 0.25, 0.50, 0.75, and 1.00 is marked with a star (*). Tables 3 and 4 also include the results for the union operation and for forward-filtered subsets of the union operation of cardinality 20 (filt20) and 50 (filt50). Space restrictions prevent us from presenting similar tables for running times and δ values for the case of a single underlying utility function. The maximum ranges for δ values under linear utility functions equal 0.35, 1.10, and 4.00, and those under Tchebycheff utility functions equal 1.20, 7.00, and 12.00, for MOKP, MOST20, and MOST50, respectively.

As might be expected, δ values increase when the form of the DM's utility function is very different from that assumed by the algorithm, when the problem dimension is high, and when the precision of the partial information on DM preferences is low. High values of the α parameter also seem to weaken the evolutionary forces involved in the algorithm, resulting in high δ levels. Simply setting $\alpha = 0.00$ and working only with the minimum utility surplus seems to yield quite successful results for MOST problems, while α levels of 0.50 or 0.25 seem to work better for MOKP. Studying the distribution of stars in Tables 3 and 4, we find that once again $\alpha = 0.25$ and 0.50 work well for MOKP, while $\alpha = 0.00$ and 0.25 perform better for MOST problems.

Table 3. Average δ_{avg} values

Problem	W	p	α					union	filt20	filt50
			0.00	0.25	0.50	0.75	1.00			
MOKP	W1	2	*0.08	*0.08	*0.08	*0.08	0.09	0.07	0.08	0.08
		3	0.11	0.11	*0.10	*0.10	0.11	0.08	0.11	0.10
		4	0.10	0.09	*0.08	0.09	0.09	0.08	0.10	0.09
	W2	2	0.10	*0.09	0.11	0.11	0.11	0.07	0.09	0.11
		3	0.15	0.13	*0.11	*0.11	0.12	0.09	0.13	0.11
		4	0.22	*0.19	*0.19	*0.19	0.22	0.14	0.21	0.22
	W3	2	0.19	*0.15	*0.15	0.17	0.17	0.09	0.15	0.15
		3	0.21	*0.16	0.17	0.20	0.23	0.12	0.17	0.18
		4	0.43	0.35	*0.34	0.35	0.38	0.25	0.39	0.38
MOST20	W1	2	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.00	0.00
		3	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.00	0.00
		4	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.01	0.04
	W2	2	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.00	0.00
		3	*0.00	*0.00	*0.00	0.01	0.01	0.00	0.04	0.02
		4	*0.04	0.05	0.09	0.10	0.18	0.02	0.18	0.13
	W3	2	*0.00	*0.00	*0.00	0.01	0.02	0.00	0.01	0.00
		3	*0.01	0.04	0.08	0.17	0.17	0.01	0.31	0.05
		4	*0.15	0.67	0.67	0.70	0.73	0.12	1.03	0.34
MOST50	W1	2	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.00	0.00
		3	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.01	0.00
		4	*0.00	*0.00	*0.00	0.01	0.01	0.00	0.04	0.01
	W2	2	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.00	0.00
		3	*0.01	0.03	0.03	0.05	0.07	0.01	0.04	0.04
		4	*0.08	0.15	0.22	0.34	0.51	0.07	0.28	0.25
	W3	2	*0.00	0.01	*0.00	0.01	0.11	0.00	0.00	0.00
		3	*0.04	0.30	0.28	0.30	0.51	0.03	0.13	0.29
		4	*0.27	0.99	1.13	1.10	1.58	0.22	0.74	1.17

Table 4. Maximum δ_{max} values

Problem	W	p	α					union	filt20	filt50
			0.00	0.25	0.50	0.75	1.00			
MOKP	W1	2	0.22	*0.16	*0.16	0.22	0.20	0.16	0.16	0.16
		3	0.27	0.29	0.27	*0.23	0.27	0.23	0.23	0.24
		4	*0.49	*0.49	*0.49	*0.49	*0.49	0.49	0.49	0.53
	W2	2	0.56	*0.27	0.39	0.28	0.29	0.18	0.18	0.18
		3	0.59	0.49	*0.44	0.46	0.51	0.35	0.35	0.36
		4	1.39	1.13	*0.91	1.65	1.69	0.83	0.83	0.92
	W3	2	0.94	0.91	*0.53	0.59	0.64	0.35	0.35	0.35
		3	1.47	*0.71	0.77	0.73	0.89	0.51	0.51	0.60
		4	3.52	*1.73	1.77	1.81	2.21	1.10	1.10	1.21
MOST20	W1	2	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.00	0.00
		3	0.02	*0.00	*0.00	0.02	0.02	0.00	0.22	0.11
		4	*0.05	*0.05	*0.05	*0.05	*0.05	0.05	0.30	0.18
	W2	2	*0.00	*0.00	*0.00	0.01	0.29	0.00	0.01	0.00
		3	0.11	*0.06	0.16	0.28	0.36	0.06	0.51	0.23
		4	*0.65	1.00	1.34	1.53	2.10	0.49	1.90	1.10
	W3	2	*0.00	*0.00	0.21	0.17	0.64	0.00	0.20	0.00
		3	*0.26	0.97	1.47	2.19	3.08	0.22	0.82	0.57
		4	*2.04	5.48	4.53	4.85	5.67	1.43	4.07	2.49
MOST50	W1	2	*0.00	*0.00	*0.00	*0.00	*0.00	0.00	0.00	0.00
		3	*0.01	*0.01	0.02	0.02	0.05	0.01	0.07	0.03
		4	*0.07	*0.07	*0.07	*0.07	*0.07	0.07	0.16	0.07
	W2	2	*0.00	0.04	0.12	0.03	0.03	0.00	0.01	0.00
		3	*0.13	0.69	0.69	0.58	0.70	0.11	0.28	0.13
		4	*0.79	1.63	2.89	3.46	3.22	0.58	1.22	0.94
	W3	2	*0.01	0.24	0.19	0.33	1.26	0.01	0.03	0.01
		3	*0.52	3.04	3.34	3.31	4.15	0.30	0.96	0.40
		4	*4.06	8.20	6.54	6.57	8.85	1.81	3.54	2.22

When the union is taken of the different versions, average δ_{avg} values remain below 0.25 %, maximum δ_{max} values below 1.81 %, and average δ values below 0.32 % with a linear underlying utility function and below 2.14 % with a Tchebycheff underlying utility function for all problems. Even if the DM does not wish to search the entire set of solutions resulting from the

union operation, the results remain encouraging for forward filtered subsets. When the DM is presented a set of at most 20 solutions, if s/he is able to choose the best therein, average δ and δ_{avg} values are below 1.2 % and δ_{max} values are below 4.5 % in all cases with linear underlying utility functions. When the sample size increases to 50, the corresponding figures fall to 1.0 % and 2.5 %. When the underlying utility function of the DM sharply differs from what is assumed by the algorithm, however, average δ values can go up to 7.0 % for 20-solution samples and 3.5 % for 50-solution samples, so conducting an interactive search over the entire union set may be worthwhile.

Computation durations, like δ values, tend to increase when there are increases in size of the problem, the dimension of the problem, or the size of the restricted weight space. We find that extreme levels of 0.00 and 1.00 for α result in shorter run times, conceivably because they simplify the LPs that must be solved for favorable weight computations. Maximum values for a 233mHz Pentium II processor equal 400 seconds for MOKP and MOST20, and 600 seconds for MOST20. If we ignore the extreme values for $\alpha = 0.75$, average durations are under 7 minutes for all cases. Combining and if necessary filtering the results of the five versions takes less than 40 minutes for all problems.

Finally, the simulation results illustrate the focusing capability of the EMAPS fitness function when one compares the union sets obtained for the same problem under different degrees of precision. Figure 3 below displays typical results for one 2-dimensional and one 3-dimensional problem. It is quite clear in the figure that as the size of the restricted weight space W increases, the cardinality and spread of the union set increases to approximate a wider range of preference-nondominated solutions.

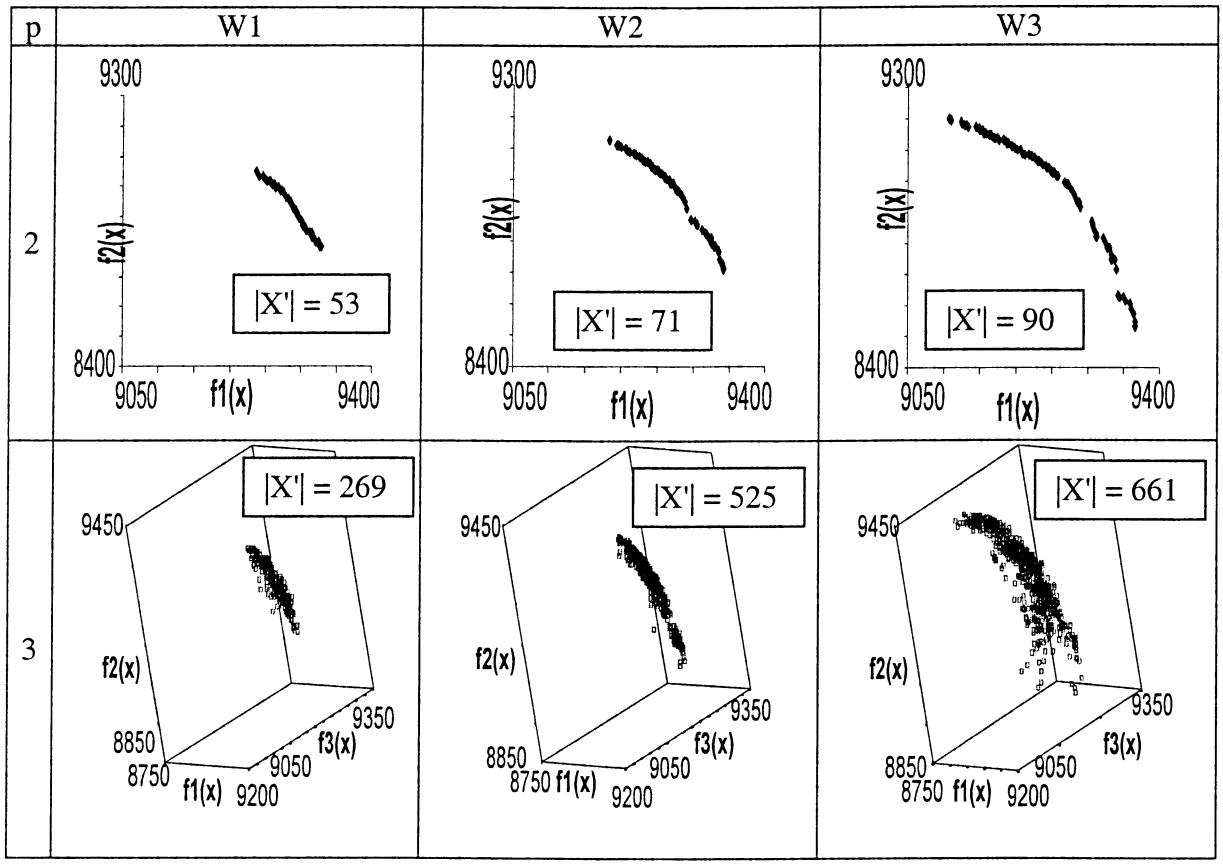

Figure 3. Illustration of the effect of precision on EMAPS

5. Conclusions

The simulation runs discussed above indicate that EMAPS succeeds in finding a high-quality approximation of the DM's preference-nondominated set in all cases. The effect of the parameter α, controlling the balance in the fitness function of the minimum and the average utility surplus, seems to follow no fixed pattern. Further research on this effect remains of interest, as does the design of a dynamic metaheuristic where α is progressively adjusted. In our simulations, α values of 0.00, 0.25, and 0.50 appeared to yield good results. The final populations reported under different levels of α were distinct but mostly of equally high quality. When the problem

size or dimension is large, or the partial information available on the DM's preferences is not very precise, the set X^{PDM} to be approximated grows quite large and running EMAPS under several different levels of α and combining the results emerges as a good alternative to increasing the maximum number of crossovers or the maximum population cardinality allowed, especially if the latter is not feasible with the given computational resources. When the resulting set from the union operation is filtered to obtain a smaller subset of disperse solutions, a representative sample of high quality solutions may be obtained.

The fitness function designed for EMAPS makes it possible for the population to move in all three of the desired evolutionary directions simultaneously without requiring the use of niching, fitness sharing, or other mechanisms traditionally employed to guarantee a good spread of the population over the efficient frontier. When a positive value is used for the parameter α, convex dominated solutions in otherwise underrepresented regions of the set of preference-nondominated solutions may have higher fitness than supported efficient solutions in crowded regions. This may improve the metaheuristic's performance for nonlinear utility functions, especially when a small population is used.

Comparisons of the algorithm with others in the literature are desirable, and may be carried out for the no information case where the entire efficient frontier is being approximated. The method can easily be customized to solve different MOCO problems as well as other multiobjective problems that cannot be solved analytically. The evolution operators and base heuristics required for customization may in many cases be readily adapted from the vast catalog of evolutionary algorithms that have been developed in recent years to address single-objective combinatorial optimization problems.

Acknowledgement

Selcen Pamuk was supported by the Scientific and Technical Research Council of Turkey (TÜBİTAK) during part of the research outlined in this paper.

References

Athanassopoulos, A.D. and V.V. Podinovski (1997), "Dominance and Potential Optimality in Multiple Criteria Decision Analysis with Imprecise Information," *Journal of the Operational Research Society*, Vol.48. pp.142-150.

Beasley, J.E. and P.C. Chu (1998), "A Genetic Algorithm for the Multidimensional Knapsack Problem," *Journal of Heuristics*, Vol.4, pp.63-86

Carrizosa, E., E. Conde, F.R. Fernandez, and J, Puerto (1995), "Multi-criteria Analysis with Partial Information about the Weighting Coefficients," *European Journal of Operational Research*, Vol.81, pp.291-301.

Czyzak, P. and A. Jaszkiewicz (1998), "Pareto Simulated Annealing - A Metaheuristic Technique for Multiple-Objective Combinatorial Optimization," *Journal of Multi Criteria Decision Analysis* Vol.7, pp.34-47.

Ehrgott, M. and X. Gandibleux (2000), "A Survey and Annotated Bibliography of Multicriteria Combinatorial Optimization". *OR Spektrum*, Vol.22, pp.425-460.

Gandibleux, X, N. Mezdaoui, and A. Freville (1996), "A Tabu Search Procedure to Solve Multiobjective Combinatorial Optimization Problems," in *Proceedings of MOPGP'96*, R. Caballero and R. Steuer (eds.), Springer-Verlag .

Hamacher, H.W. and G. Ruhe (1994), "On Spanning Tree Problems with Multiple Objectives," *Annals of Operations Research,* Vol.52, pp.209-230.

Hazen, G.B. (1986), "Partial Information, Dominance, and Potential Optimality in Multiattribute Utility Theory," *Operations Research*, Vol.34, No.2, pp.296-310.

Kirkwood, C.W. and R.K. Sarin (1985), "Ranking with Partial Information: A Method and an Application," *Operations Research,* Vol.33, pp.38-48.

Malakooti, B, L. Duckstein, and A. Ravindan (1994), "Screening Discrete Alternatives with Imprecisely Assessed Additive Multi-Attribute Functions," *Applied Mathematics and Computation*, Vol.60, pp.69-87.

Martello, S. and P. Toth (1990), *Knapsack Problems: Algorithms and Computer Implementations,* John Wiley and Sons, Chichester

Pamuk, S. and M. Köksalan (2001), "An Interactive Evolutionary Metaheuristic for Multiobjective Combinatorial Optimization," DCMME Working Paper Series 01-01-1, Purdue University.

Reeves, C.R., ed., (1993), *Modern Heuristic Techniques for Combinatorial Problems*, John Wiley and Sons, Ltd., New York.

Steuer, R.E. (1986), *Multiple criteria optimization: Theory, Computation, and Application*, John Wiley and Sons, New York.

Stewart, T. J. (1992), "A Critical Survey on the Status of Multiple Criteria Decision Making Theory and Practice," *OMEGA International Journal for Management Science* Vol.20, pp.569-586.

Takeda, E. and J. Satoh (2000), "A Data Envelopment Analysis Approach to Multicriteria Decision Problems with Incomplete Information," *Computers and Mathematics with Applications,* Vol.39, pp.81-90.

Teghem, J., D. Tuyttens, and E.L. Ulungu (2000), "An Interactive Heuristic Method for Multiobjective Combinatorial Optimization," *Computers and Operations Research*, Vol.27, pp.621-634.

Ulungu E.L. and J. Teghem (1994), "Multi-Objective Combinatorial Optimization Problems: A Survey," *Journal of Multi Criteria Decision Analysis* Vol.3, pp.83-104.

Weber, M. (1987), "Decision Making with Incomplete Information," *European Journal of Operational Research,* Vol.28, pp.44-57.

Zitzler, E. and L. Thiele (1999), "Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach," *IEEE Transactions on Evolutionary Computation*, Vol.3, No.4, pp.257-271.

ADDITIONAL INSTITUTE PAPERS AVAILABLE FROM THE KRANNERT GRADUATE SCHOOL OF MANAGEMENT

-1995-

1069 Sugato Chakravarty and John J. McConnell, AN ANAYLSIS OF PRICES, BID/ASK SPREADS, AND BID AND ASK DEPTH SURROUNDING IVAN BOESKY'S ILLEGAL TRADING IN CARNATION'S STOCK.

1070 John J. McConnell and Henri Servaes, EQUITY OWENERSHIP AND THE TWO FACES OF DEBT.

1071 Kenneth J. Matheny, REAL EFFECTS OF MONETARY POLICY IN A 'NEOCLASSICAL' MODEL: THE CASE OF INTEREST RATE TARGETING.

1072 Julie Hunsaker and Dan Kovenock, THE PATTERN OF EXIT FROM DECLINING INDUSTRIES.

1073 Kessan Joseph, Manohar U. Kalwani, THE IMPACT OF ENVIRONMENTAL UNCERTAINTY ON THE DESIGN OF SALESFORCE COMPENSATION PLANS.

1074 K. Tomak, A NOTE ON THE GOLDFELD QUANDT TEST

1075 Alok R. Chaturvedi, SIMDS: A SIMULATION ENVIRONMENT FOR THE DESIGN OF DISTRIBUTED DATABASE SYSTEMS

1076 Dan Kovenock and Suddhasatwa Roy, FREE RIDING IN NON-COOPERATIVE ENTRY DETERRENCE WITH DIFFERENTIATED PRODUCTS

1077 Kenneth Matheny, THE MACROECONOMICS OF SELF-FULFILLING PROPHECIES

1078 Paul Alsemgeest, Charles Noussair and Mark Olson, EXPERIMENTAL COMPARISONS OF AUCTIONS UNDER SINGLE-AND MULTI-UNIT DEMAND

1079 Dan Kovenock, Casper D de Vries, FIAT EXCHANGE IN FINITE ECONOMIES

1080 Dan Kovenock, Suddhasatwa Roy, DYNAMIC CAPACITY CHOICE IN A BERTRAND-EDGEWORTH FRAMEWORK

1081 Burak Kazaz, Canan Sepil, PROJECT SCHEDULING WITH DISCOUNTED CASH FLOWS AND PROGRESS PAYMENTS

-1996-

1082 Murat Koksalan, Oya Rizi, A VISUAL INTRACTIVE APPROACH FOR MULTIPLE CRITERIA DECISION MAKING WITH MONOTONE UTILITY FUNCTIONS

1083 Janet S. Netz, John D. Haveman, ALL IN THE FAMILY: FAMILY, INCOME, AND LABOR FORCE ATTACHMENT

1084 Keith V. Smith, ASSET ALLOCATION AND INVESTMENT HORIZON

1085 Arnold C. Cooper and Catherine M. Daily, ENTREPRENEURIAL TEAMS

1086 Alok R. Chaturvedi and Samir Gupta, SCHEDULING OF TRANSACTIONS IN A REAL-TIME DISTRIBUTED TRANSACTION PROCESSING SYSTEMS: SCALEABILITY AND NETWORKING ISSUES

1087 Gordon P. Wright, N. Dan Worobetz, Myong Kang, Radha V. Mookerjee and Radha Chandrasekharan, OR/SM: A PROTOTYPE INTEGRATED MODELING ENVIRONMENT BASED ON STRUCTURED MODELING

1088 Myong Kang, Gordon P. Wright, Radha Chandrasekharan, Radha Mookerjee and N. Dan Worobetz, THE DESIGN AND IMPLEMENTATION OF OR/SM: A PROTOTYPE INTEGRATED MODELING ENVIRONMENT

1089 Thomas H. Brush and Philip Bromiley, WHAT DOES A SMALL CORPORATE EFFECT MEAN? A VARIANCE COMPONENTS SIMULATION OF CORPORATE AND BUSINESS EFFECTS

1090 Kenneth J. Matheny, NON-NEUTRAL RESPONSES TO MONEY SUPPLY SHOCKS WHEN CONSUMPTION AND LEISURE ARE PARETO SUBSTITUTES

1091 Kenneth J. Matheny, MONEY, HUMAN CAPITAL, AND BUSINESS CYCLES: A MODERN PHILLIPS CURVE-STYLE TRADEOFF

1092 Kenneth J. Matheny, OUTPUT TARGETING AND AN ARGUMENT FOR STABILIZATION POLICIES

1093 Kenneth J. Matheny, THE RELEVANCE OF OPEN MARKET OPERATIONS AS A MONETARY POLICY TOOL

-1997-

1094 James C. Moore, William Novshek and Peter Lee U, ON THE VOLUNTARY PROVISION OF PUBLIC GOODS

1095 Michael R. Baye, Dan Kovenock and Casper G. deVries, THE INCIDENCE OF OVERDISSIPATION IN RENT-SEEKING CONTESTS

1096 William Novshek and Lynda Thoman, CAPACITY CHOICE AND DUOPOLY INCENTIVES FOR INFORMATION SHARING

1097 Vidyanand Choudhary, Kerem Tomak and Alok Chaturvedi, ECONOMIC BENEFITS OF RENTING SOFTWARE

1098 Jeongwen Chiang and William T. Robinson, DO MARKET PIONEERS MAINTAIN THEIR INNOVATIVE SPARK OVER TIME?

1099 Glenn Hueckel, LABOR COMMAND IN *THE WEALTH OF NATIONS*: A SEARCH FOR "SYSTEM"

1100 Glenn Hueckel, SMITH'S UNIFORM "TOIL AND TROUBLE": A "VAIN SUBTLETY"?

1101 Thomas H. Brush and Philip Bromiley, WHAT DOES A SMALL CORPORATE EFFECT MEAN? A VARIANCE COMPONENTS SIMULATION OF CORPORATE AND BUSINESS EFFECTS

1102 Thomas Brush, Catherine Maritan and Aneel Karnani, MANAGING A NETWORK OF PLANTS WITHIN MULTINATIONAL FIRMS

1103 Sam Hariharan and Thomas H. Brush, RESOURCES AND THE SCALE OF ENTRY CHOICE: THE COMPETITIVE ADVANTAGE OF ESTABLISHED FIRMS?

1104 Thomas H. Brush, Philip Bromiley and Margaretha Hendrickx, THE RELATIVE INFLUENCE OF INDUSTRY AND CORPORATION ON BUSINESS SEGMENT PERFORMANCE: AN ALTERNATIVE ESTIMATE

1105 Thomas Brush, Catherine Maritan and Aneel Karnani, PLANT ROLES IN THE MANAGEMENT OF MULTINATIONAL MANUFACTURING FIRMS

1106 Thomas H. Brush, Catherine Maritan and Aneel Karnani, THE PLANT LOCATION DECISION IN MULTINATIONAL MANUFACTURING FIRMS: AN EMPIRICAL ANALYSIS OF INTERNATIONAL BUSINESS AND MANUFACTURING STRATEGY PERSPECTIVES

1107 Piyush Kumar, Manohar U. Kalwani and Maqbool Dada, THE IMPACT OF WAITING TIME GUARANTEES ON CUSTOMERS' WAITING EXPERIENCES

1108 Thomas H. Brush, Philip Bromiley and Margaretha Hendrickx, THE FREE CASH FLOW HYPOTHESIS FOR SALES GROWTH AND FIRM PERFORMANCE

1109 Keith V. Smith, PORTFOLIO ANALYSIS OF BROKERAGE FIRM RECOMMENDATIONS

- 1998 -

1110 Charles Noussair, Kenneth Matheny, and Mark Olson, AN EXPERIMENTAL STUDY OF DECISIONS IN DYNAMIC OPTIMIZATION PROBLEMS

1111 Jerry G. Thursby and Sukanya Kemp, AN ANALYSIS OF PRODUCTIVE EFFICIENCY OF UNIVERSITY COMMERCIALIZATION ACTIVITIES

1112 John J. McConnell and Sunil Wahal, DO INSTITUTIONAL INVESTORS EXACERBATE MANAGERIAL MYOPIA?

1113 John J. McConnell, Mehmet Ozbilgin and Sunil Wahal, SPINOFFS, EX ANTE

1114 Sugato Chakravarty and John J. McConnell, DOES INSIDER TRADING REALLY MOVE STOCK PRICES?

1115 William T. Robinson and Sungwook Min, IS THE FIRST TO MARKET THE FIRST TO FAIL?: EMPIRICAL EVIDENCE FOR MANUFACTURING BUSINESSES

1116 Margaretha Hendrickx, WHAT CAN MANAGEMENT RESEARCHERS LEARN FROM DONALD CAMPBELL, THE PHILOSOPHER? AN EXERCISE IN PHILOSOPHICAL HERMENEUTICS

1117 Thomas H. Brush, Philip Bromiley and Margaretha Hendrickx, THE FREE CASH FLOW HYPOTHESIS FOR SALES GROWTH AND FIRM PERFORMANCE

1118 Thomas H. Brush, Constance R. James and Philip Bromiley, COMPARING ALTERNATIVE METHODS TO ESTIMATE CORPORATE AND INDUSTRY EFFECTS

1119 Charles Noussair, Stéphane Robin and Bernard Ruffieux, BUBBLES AND ANTI-CRASHES IN LABORATORY ASSET MARKETS WITH CONSTANT FUNDAMENTAL VALUES

1120 Vivian Lei, Charles N. Noussair and Charles R. Plott, NON-SPECULATIVE BUBBLES IN EXPERIMENTAL ASSET MARKETS: LACK OF COMMON KNOWLEDGE OF RATIONALITY VS. ACTUAL IRRATIONALITY

-1999-

1121 Kent D. Miller and Timothy B. Folta, ENTRY TIMING AND OPTION VALUE

1122 Glenn Hueckel, THE LABOR "EMBODIED" IN SMITH'S LABOR-COMMANDED MEASURE: A "RATIONALLY RECONSTRUCTED" LEGEND

1123 Timothy B. Folta and David A. Foote, TEMPORARY EMPLOYEES AS REAL OPTIONS

1124 Gabriele Camera, DIRTY MONEY

1125 Wilfred Amaldoss, Robert J. Meyer, Jagmohan S. Raju, and Amnon Rapoport, COLLABORATING TO COMPETE: A GAME-THEORETIC MODEL AND EXPERIMENTAL INVESTIGATION OF THE EFFECT OF PROFIT-SHARING ARRANGEMENT AND TYPE OF ALLIANCE ON RESOURCE-COMMITMENT DECISIONS

1126 Wilfred Amaldoss, Robert J. Meyer, Jagmohan S. Raju, and Amnon Rapoport, APPENDICES FOR COLLABORATING TO COMPETE: A GAME-THEORETIC MODEL AND EXPERIMENTAL INVESTIGATION OF THE EFFECT OF PROFIT-SHARING ARRANGEMENT AND TYPE OF ALLIANCE ON RESOURCE-COMMITMENT DECISIONS

-2000-

1127 Sugato Chakravarty and Kai Li, AN ANALYSIS OF OWN ACCOUNT TRADING BY DUAL TRADERS IN FUTURES MARKETS: A BAYESIAN APPROACH

1128 Sugato Chakravarty, STEALTH TRADING: THE NEXT GENERATION

1129 S.G. Badrinath and Sugato Chakravarty, ARE ANALYST RECOMMENDATIONS INFORMATIVE?

1130 Sugato Chakravarty and Asani Sarkar, THE DETERMINANTS OF LIQUIDITY IN U.S. CORPORATE, MUNICIPAL AND TREASURY BOND MARKETS

1131 Vivian Lei and Charles Noussair, AN EXPERIMENTAL TEST OF AN OPTIMAL GROWTH MODEL

1132 Paul Healy and Charles Noussair, BIDDING BEHAVIOR IN THE PRICE IS RIGHT GAME: AN EXPERIMENTAL STUDY

1133 Kent D. Miller and Zur Shapira, BEHAVIORAL OPTION THEORY: FOUNDATIONS AND EVIDENCE

1134 Kent D. Miller, KNOWLEDGE INVENTORIES AND MANAGERIAL MYOPIA

1135 Gabriele Camera, Charles Noussair, and Steven Tucker, RATE-OF-RETURN DOMINANCE AND EFFICIENCY IN AN EXPERIMENTAL ECONOMY

1136 Timothy B. Folta, Jay J. Janney, SIGNALING FOR RESOURCE ACQUISITION: PRIVATE EQUITY PLACEMENTS BY TECHNOLOGY FIRMS

1137 Michael R. Baye, Dan Kovenock, Casper G. de Vries, COMPARATIVE ANALYSIS OF LITIGATION SYSTEMS: AN AUCTION-THEORETIC APPROACH

1138 Sugato Chakravarty, Asani Sarkar, DO DIFFERENCES IN TRANSPARENCY AFFECT TRADING COSTS? EVIDENCE FROM U.S. CORPORATE, MUNICIPAL AND TREASURY BOND MARKETS

-2001-

1139 Charles Noussair, Stephane Robin, Bernard Ruffieux, GENETICALLY MODIFIED ORGANISMS IN THE FOOD SUPPLY: PUBLIC OPINION VS CONSUMER BEHAVIOR

1140 Gabriele Camera, SEARCH, DEALERS, AND THE TERMS OF TRADE

1141 David Masclet, Charles Noussair, Steven Tucker, Marie-Claire Villeval, MONETARY AND NON-MONETARY PUNISHMENT IN THE VOLUNTARY CONTRIBUTIONS MECHANISM

1142 Charles Noussair, Stephane Robin, Bernard Ruffieux, DO CONSUMERS NOT CARE ABOUT BIOTECH FOODS OR DO THEY JUST NOT READ THE LABELS

1143 Timothy B. Folta, Douglas R. Johnson, Jonathan O'Brien, UNCERTAINTY AND THE LIKELIHOOD OF ENTRY: AN EMPIRICAL ASSESSMENT OF THE MODERATING ROLE OF IRREVERSIBILITY

1144 Alok Chaturvedi, Mukul Gupta, Sameer Gupta, ISSUES IN SERVER FARM DESIGN FOR REAL TIME E-COMMERCE TRANSACTIONS

1145 William Blankenau, Gabriele Camera, PRODUCTIVE EDUCATION OR A MARKETABLE DEGREE?