

ASSET REPLACEMENT CONSIDERING ENVIRONMENTAL AND ECONOMIC OBJECTIVES

by
Darby E. Grande

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in The University of Michigan
2004

Doctoral Committee:

Professor James C. Bean, Chair
Assistant Professor Sebastian Fixson
Associate Professor Gregory A. Keoleian
Assistant Professor Dushyant Sharma

© Darby E. Grande 2004
All Rights Reserved

ACKNOWLEDGEMENTS

I am truly humbled by the many contributions that led to the realization of this work. The first thanks are due to my advisor, James Bean, whose breadth of knowledge, insightful mentoring, and example of leadership have kept me working toward this goal. I also thank my committee members Sebastian Fixson, Greg Keoleian and Dushyant Sharma for their valuable help and advice. This work would not have been possible without the Life Cycle Optimization of Vehicle Replacement team, including Greg Keoleian, Jonathan Bulkley, Hyung Chul Kim, James Bean, Marc Ross, and research partners at General Motors and the EPA. Additionally, I appreciate David Spitzley's work with the cost modeling. I would like to acknowledge the financial support of the National Science Foundation under Grant No. BES-9985625 as well as the Horace H. Rackham School of Graduate Studies and the Department of Industrial and Operations Engineering at the University of Michigan.

I am extremely grateful for the support I was given by the IOE department as a whole. I owe much to Larry Seiford, our Department Chair, who has been a constant source of advice on networking, teaching, time management, and leadership. I am indebted to the wonderful staff who lessened the burden of all things administrative and technical. In the faculty I have found a number of role models, and I am especially grateful to the exciting cohort of junior faculty who shared their experience and advice. I appreciate the many graduate student colleagues that have shared time, knowledge and encouragement through the many stages of my time here. These

friendships and those of many outside the IOE community have been an incredible source of strength and joy.

Finally, none of this would have been possible without the blessing of such a supportive family. To my dear parents, whose example in the lifelong pursuit of knowledge will be with me always, I am unable express in words my gratitude for your many means of love and support. For Erik's creativity, Mimi's thoughtful phone calls, and all of my extended family, I give heartfelt thanks.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
LIST OF FIGURES	vi
LIST OF TABLES	viii
CHAPTER	
I. Introduction	1
II. Vehicle Replacement	7
2.1 Assumptions	8
2.2 Data	9
2.3 Single-vehicle Replacement Model	11
2.3.1 Integer Program Formulation	12
2.3.2 Dynamic Program Formulation	13
2.4 Single-objective Results	15
2.4.1 Refinement of Objectives	16
2.4.2 Objective Trade-offs	17
III. Multiobjective Metaheuristics	18
3.1 Single-objective Heuristic Algorithms	20
3.1.1 Genetic Algorithms	20
3.1.2 Tabu Search	23
3.2 Multiobjective Heuristic Algorithms	24
3.2.1 Strength Pareto Evolutionary Algorithm (SPEA/M)	26
3.2.2 Nondominated Sorting Genetic Algorithm II (NSGA-II/M)	27
3.2.3 Multiobjective Tabu Search (MOTS/M)	29
3.3 Implementation	32
3.3.1 Novel Algorithm Details	33
3.3.2 Verifying Pareto Optimality	34
3.3.3 Algorithm Performance	36
3.3.4 Algorithm Comparison	40
3.4 Remarks and Conclusions	41
IV. Multivehicle Study	44
4.1 Policy Analysis of Complex Systems	45
4.2 Model	46
4.3 Data	47
4.3.1 Household Profile	48

4.3.2	Baseline Consumer Behavior	48
4.3.3	Policy Scenarios	50
4.4	Results	52
4.4.1	Fuel Tax	53
4.4.2	Hybrid Only	54
4.4.3	Fuel Tax + Hybrid	55
4.4.4	Multiple Challengers	55
4.4.5	Alternative Consumer Behavior	57
4.5	Discussion	60
V. Conclusions		71
BIBLIOGRAPHY		74

LIST OF FIGURES

Figure

2.1	Magnitude of Trade-offs	17
3.1	Pareto-optimal fronts sorting	28
3.2	Pareto-optimal solutions after 100 SPEA/M generations	38
3.3	Pareto-optimal solutions after 100 NSGA-II/M generations	42
3.4	MOTS/M solutions with Pareto-optimal solutions for comparison	43
3.5	Pareto-optimal solutions as generations increase	43
4.1	Policy Decision Support System Framework	47
4.2	Average Annual Cost: Fuel Tax vs. Baseline scenarios	54
4.3	Average Annual $NMHC + NO_x$ emissions: Fuel Tax vs. Baseline scenarios	55
4.4	Average Annual Cost: Hybrid Only vs. Baseline scenarios	57
4.5	Average Annual $NMHC + NO_x$ emissions: Hybrid Only vs. Baseline scenarios	58
4.6	Average Annual Cost: Fuel Tax + Hybrid vs. Baseline scenarios	59
4.7	Average Annual $NMHC + NO_x$ emissions: Fuel Tax + Hybrid vs. Baseline scenarios	60
4.8	Average Annual Cost: Multiple Challengers vs. Baseline scenarios	61
4.9	Average Annual $NMHC + NO_x$ emissions: Multiple Challengers vs. Baseline scenarios	63
4.10	Average Annual Cost: Alternative Consumer Behavior 1 vs. Baseline scenarios	64
4.11	Average Annual $NMHC + NO_x$ emissions: Alternative Consumer Behavior 1 vs. Baseline scenarios	65
4.12	Average Annual Cost: Alternative Consumer Behavior 2 (20%) vs. Baseline scenarios	66
4.13	Average Annual $NMHC + NO_x$ emissions: Alternative Consumer Behavior 2 (20%) vs. Baseline scenarios	67
4.14	Average Annual Cost: Alternative Consumer Behavior 2 (50%) vs. Baseline scenarios	67

4.15	Average Annual <i>NMHC</i> + <i>NO_x</i> emissions: Alternative Consumer Behavior 2 (50%) vs. Baseline scenarios	68
4.16	Average Annual Cost: Alternative Consumer Behavior 3 vs. Baseline scenarios . . .	68
4.17	Average Annual <i>NMHC</i> + <i>NO_x</i> emissions: Alternative Consumer Behavior 3 vs. Baseline scenarios	69
4.18	Cumulative Economic Cost	69
4.19	Cumulative <i>NMHC</i> + <i>NO_x</i> Emissions	70
4.20	Cumulative <i>CO₂</i> Emissions	70

LIST OF TABLES

Table

2.1	Single-objective Optimal Policies (1985 - 2020)	16
3.1	SPEA/M Pseudocode	27
3.2	NSGA-II/M Pseudocode	30
3.3	MOTS/M Pseudocode	32
3.4	Pareto-optimal solutions found in 100 generations of the SPEA/M	36
3.5	Pareto-optimal solutions found in 100 generations of the NSGA-II/M	39
3.6	Outline of the three algorithms	40
3.7	Variance within the trials as generations increase	41
4.1	Baseline Consumer Behavior	50
4.2	Data for Hybrid Vehicle Profile	53
4.3	Consumer behavior from Hybrid Only scenario	56
4.4	Consumer behavior from Fuel Tax + Hybrid scenario	56
4.5	Consumer behavior from Alternative Consumer Behavior 1 scenario	62
4.6	Cumulative burdens over the 1995 - 2050 horizon	62
4.7	Population-wide age averages for each scenario	63

CHAPTER I

Introduction

The U.S. Environmental Protection Agency estimates that in 1996 there were more than 120 million light-duty vehicles on the road in the United States [36]. According to the U.S. Department of Transportation, individuals averaged four trips per day in 2001, 87% of which were taken in personal vehicles [60]. The destinations for these trips include family and personal reasons such as shopping and errands, social and recreational trips for vacation or to see friends, and work-related commuter travel [60]. As personal vehicles are involved in so many aspects of daily life in the United States, the sustainability of this mode of travel is of national interest.

Transportation sources accounted for 67% of the petroleum consumption in the United States in 1999, and net imports fulfilled 49.6% of the total 1999 U.S. petroleum consumption [15]. Additionally, the transportation system causes adverse impacts on the environment which include emissions of criteria air pollutants and greenhouse gases, habitat and land use, harm to water quality and aquatic resources, hazardous materials incidents, noise, and creation of solid waste [36]. These energy and environmental concerns contribute to the complexity of analysis and decision-making within the transportation sector.

Consider the problem of determining when, and with what, to replace a passenger

vehicle. This decision will impact economics, emissions and energy use. The best decision for economics may not be a good decision for emissions. How do we trade-off these various objectives to find the “right” decision?

Asset replacement problems that minimize the cost (or maximize the profit) to the decision-maker have been studied for many years (eg. [59], [9], [55], [38], [37]). There is, however, a lack of literature addressing replacement problems in which conflicting objectives must be considered. The study of multiobjective optimization is motivated by the fact that “many decision problems have to be placed in a much broader framework than covered by the interest of the decision-maker” [47]. The literature concerning multiobjective analysis therefore attempts to aid decision-makers who must consider more than one criterion at a time.

Multiobjective analysis has its origins in microeconomics theory from the late nineteenth century [58]. Optimization literature on the subject has been growing since the late 1960s [10]. For the purposes of this discussion, we define a multiobjective optimization problem as a mathematical program with objective function that is a vector of attributes to be maximized or minimized [10]. An important difference encountered when optimizing over several objectives is that one is not likely to find a single solution that is best (a global optimum), as is usually the result with single-objective optimization [57]. The aim, therefore, is to present the decision-maker with a set of good solutions from which he or she may choose according to the specific agenda of the particular decision-maker. Because the existence of a single solution attaining the desired extreme values of all components of the objective is rare, we wish to trade-off alternate strong solutions. That is, within the set of feasible solutions, we seek a solution that is “good” in all objectives, conceding that it is unlikely to be “best” for any objective.

Traditional approaches to multiobjective optimization have included some techniques in which the decision-maker must specify upfront the relative importance of each objective under consideration. An example of this is the weighting method, in which a single objective is created by taking the inner product of the objective vector with a vector of predetermined weights [10]. Another traditional approach is an adaptation of goal programming in which the decision-maker specifies acceptable ranges for the objectives, and a single solution to minimize deviation from the goals is found [52]. Similar to this approach is the ε -constraint method [10] in which all objectives but one are constrained, and the final objective is minimized. By varying the constraints on the objectives and alternating which objective is minimized, the ε -constraint method can give a decision-maker a good indication of the available solutions. The many iterations required, and the fine mesh of constraints which must be employed in order that solutions are not missed, can make implementation of such methods extremely tedious.

Due to the fact that traditional approaches to multiobjective optimization are limited to discovering one solution at a time, interest in population-based evolutionary algorithms, which may find multiple new solutions each generation, has grown quickly in the past 20 years [17]. Because we are able to find multiple solutions more easily, the first stages of the optimization are not limited to determining a single, “best” solution. Instead we would like to construct the set of Pareto-optimal solutions, that is, the set of solutions for which the improvement in one objective requires degrading the value of another objective. From this set, a decision-maker may glean valuable trade-off information, eventually leading him or her to the desired “right” solution for the specific circumstances.

For two solutions, x_1 and x_2 , to a minimization problem with vector objective

function $f(x) = [f_1(x), \dots, f_K(x)]$, we say $f(x_1)$ *dominates* $f(x_2)$ if and only if $f_i(x_1) \leq f_i(x_2)$ for all $i \in \{1, \dots, K\}$ and there is at least one $i \in \{1, \dots, K\}$ for which $f_i(x_1) < f_i(x_2)$. In this case, x_2 is *inferior* to x_1 . For a given set, Ω , of solutions to the problem, any $x_j \in \Omega$ which is not inferior to any other solution in Ω is referred to as *noninferior* with respect to Ω , and its objective vector $f(x_j)$ is locally *nondominated*. We reserve the term ‘‘Pareto-optimal’’ for those solutions which are noninferior with respect to the entire decision variable space of the problem under consideration. We note that the aim with each algorithm’s implementation may be to construct a representative set of Pareto-optimal solutions rather than the entire Pareto-optimal set.

Construction of multiobjective, metaheuristic algorithms to find a set of Pareto-optimal solutions allows us to gain insight into tradeoffs over the life of a single vehicle. This can help inform the decisions of individual consumers or automobile manufacturers concerned with determining a vehicle’s design life. A third category of stakeholders that influences the composition of the population of vehicles on the road at any time is the national policymakers. These decision-makers implement policies that may have broad impacts, such as the CAFE (Corporate Average Fuel Economy) standards restricting the profile of the vehicles for sale each year, fuel costs and taxes, infrastructure development for alternative fuels and incentives for consumers to purchase vehicles with new technology.

In order to provide information regarding the impacts of such policy changes, we must broaden our analysis from the single vehicle to a population-level perspective. Walker [63] defines public policy analysis as ‘‘a process that generates information on the consequences that would follow the adoption of various policies,’’ and this is the information we would like to present to transportation-related decision-makers

at the national policy level. Our approach to developing an understanding of these consequences is to construct a model in which we aggregate the decisions of a variety of vehicle owners. These consumers are influenced by each other, the choices made available to them by the automobile manufacturers, public policy mandates, and their physical and economic environments. The model aggregating these consumers is therefore quite complex, and we look to the discipline of agent-based models for experience in constructing the decision-support system for policymakers.

Multiagent models have been used in a wide variety of large-scale applications. Elliot and Kiel [23] propose an agent-based model for understanding terrorist organizations, Bui [7] describes a multiagent approach to telemedicine, and Kwon and Sadeh [49] develop a multiagent architecture to enable a new understanding of comparative shopping. There are limitations of the large-scope, complex models which are somewhat reflected in Sterman's statement that "When we point to outside shocks and side effects to excuse the failure of our policies, we think we are describing a capricious and unpredictable reality. In fact, we are highlighting the limitations of our mental models" [39]. Models involving environmental impacts and reactions have, however, received encouraging attention from the multiagent, complex systems perspective (see [13], [65], [20], [34], and [25]), and we work take these concerns into consideration as we approach our multivehicle, multiobjective replacement problem.

In Chapter II we discuss the specific vehicle replacement problem analyzed as part of this work, and we describe the data utilized here. In Chapter III we introduce multiobjective heuristic algorithms, including a detailed description of the three implemented for this study. We discuss the outcomes of the algorithm implementations and consider additional analysis of the algorithms. In Chapter IV we expand our analysis to the nation-wide population of vehicles, and we consider the impact

on this population of changes in national policy, such as an increase in fuel tax or emissions regulations. From these anticipated impacts, we examine the change in the environmental and economic burden of the vehicle population due to such policy changes. Finally, we draw conclusions and discuss future work in Chapter V.

CHAPTER II

Vehicle Replacement

Vehicle service lifetime decisions of automobile manufacturers, policymakers and environmentally concerned consumers are influenced by the energy use and harmful emissions associated with owning and driving a vehicle. We begin our vehicle replacement analysis by examining the life of a single vehicle, and this serves as the foundation for the nationwide model discussed in Chapter IV. Using data gathered from a variety of industrial and governmental sources, we use the Industrial Ecology methodology of Life Cycle Assessment (LCA) to construct profiles of a generic, midsize vehicle's energy use, cost and emissions produced. Upon the retirement of this vehicle, we introduce a new vehicle, whose LCA profiles are adjusted to reflect anticipated advances in technology over time, to perform the equivalent function of the first vehicle. We continue these replacements for the length of the study so that a sequence of vehicles and their optimal retirement ages is constructed.

Our ultimate aim in this single-vehicle study is to determine the service lifetimes able to best satisfy a decision-maker's concern for all of the (conflicting) objectives considered: cost, energy use and HC , CO , CO_2 , NO_x emitted. This will necessarily involve a trade-off among the objectives since, for example, newer vehicles which use less energy and produce fewer emissions require a greater economic investment

than the maintenance of older, less efficient vehicles. Because the objectives are noncommensurable, the construction of this trade-off is nontrivial and requires an understanding of the behavior of the individual objectives. Therefore, we begin with the analysis of several single objective, single-vehicle replacement problems, and then move to the multiobjective analysis discussed in Chapter III.

2.1 Assumptions

Many older vehicles, instead of being scrapped immediately upon the purchase of a replacement, are kept as secondary vehicles. These vehicles are not driven daily, yet they have not been dismantled. The sequential nature of our model requires that the challengers perform the same service as the defender so that their behaviors may be compared. In order to ensure that each vehicle under consideration in the model performs an equivalent function over its entire life, we limit our definition of vehicle life to be the number of years for which the vehicle remains a member of the population of *Primary Use Vehicles*. We characterize this population by the following:

- Each vehicle is a mid-size, 5-passenger domestic sedan.
- The vehicle is assumed to be driven 12,000 miles per year.
- Every owner maintains the vehicle according to the manufacturer's recommendations, and all maintenance is performed by a service professional.
- Only regular unleaded gasoline is used.

As long as the vehicle remains a Primary Use Vehicle, it may be owned by a single driver or a series of drivers. The environmental factors are unaffected by a change in ownership since, for our purposes, the vehicle's function never changes. Additionally,

we assume that financial transactions at the time of an ownership change are completed according to the vehicle's depreciated value so that no additional vehicle cost is created during an ownership change. Therefore the only purchase cost incurred per vehicle is the initial purchase of the vehicle when it enters the population of Primary Use Vehicles.

When a vehicle is no longer a Primary Use Vehicle, it may enter another population of vehicles with an alternative function (a family's weekend car, say) or be scrapped. For the purposes of this study, we consider that all vehicles leaving the Primary Use population are scrapped immediately since in this case the ensuing patterns of use are shown to be reduced well below the annual mileage for the households' primary vehicles [15].

2.2 Data

A collaboration with the School of Natural Resources, the Department of Physics, and the Department of Industrial and Operations Engineering at the University of Michigan, in addition to industry partners at General Motors and the U.S. Environmental Protection Agency's Vehicle Emissions Laboratory, has provided a rich data set for our study of the multiobjective asset replacement problem. For every vehicle and year considered in the model, we require several parameters in order to determine its optimal retirement age:

- Vehicle Characteristics
 - The fuel economy of each new vehicle model as well as its predicted fuel economy performance over time.
 - The assumed number of vehicle miles traveled per year.

- Economic costs:
 - The purchase price of the vehicle (assumed to reflect the cost of manufacturing the vehicle)
 - The annual cost to maintain the vehicle for each year of its life
 - The average annual price of gasoline and the fuel economy of each new vehicle model as well as its predicted fuel economy performance over time
 - The annual ownership costs (insurance, registration and taxes)
 - The value to be received by the last owner upon disposal of the vehicle

- Energy Used:
 - The energy required for manufacturing the vehicle
 - The energy used in the preparation and use of the fuel per mile driven (the precombustion and combustion of each gallon of gasoline) for each calendar year and each vehicle model
 - The energy used during vehicle service (overhead as well as manufacturing of new parts)
 - The energy used in the end of life processing of the vehicle

- Emissions (for each pollutant under consideration):
 - The quantity of emissions released during materials production and vehicle manufacturing
 - The emissions released in the preparation and use of the fuel per mile driven (the precombustion and combustion of gasoline)
 - The emissions released during maintenance and the production of parts required for maintenance

- The quantity of emissions released during the end of life processing of the vehicle.

By varying these parameters we are able to examine the impact of each on the optimal retirement ages of the vehicles. This flexibility is of particular importance when considering future vehicles. The life cycle specifications change with each technological improvement, and several significant changes in vehicle and fuel technologies are expected in the near future [45]. A model for decision-making that depends on the future must therefore consider several potential improvement schedule scenarios. A complete description of the development of the data profiles is available in [43].

2.3 Single-vehicle Replacement Model

To begin our study of the multiobjective, vehicle-replacement problem, we consider the optimization of the replacement of a single, mid-size, passenger sedan with a comparable new vehicle to perform the same service (a constant number of miles are driven per year). At the start of each year, the owner makes the decision to either buy a new vehicle or keep the existing vehicle. The solutions, therefore, are series of replacement decisions of the form “Buy a new vehicle at the start of year 1985 and keep it for 10 years; then buy a new vehicle at the start of year 1995 and keep it for 6 years, etc.”

We currently assume a deterministic model in which there is one new vehicle available for purchase each year, and replacement decisions are made only at the beginning of each year. Additionally, the time horizon is a finite number of years in which decisions are made, and all vehicles have a finite maximum service life after which they must be retired.

2.3.1 Integer Program Formulation

The following parameters and variables are used in constructing an integer program to model the replacement problem:

$$\begin{aligned}
 N &= \text{The time horizon for the study} \\
 V &= \text{The maximum service life of a single vehicle} \\
 K &= \text{The number of objectives to minimize} \\
 x_{ij} &= \begin{cases} 1 & \text{if purchase a vehicle at the beginning of year } i \\ & \text{and keep through year } j; \\ & i \in \{1, 2, \dots, N\}, j \in \{i, i+1, \dots, i+V-1\} \\ 0 & \text{otherwise} \end{cases} \\
 c_{kij} &= \text{The cost to objective } k \text{ of purchasing a vehicle in year } i \\ & \text{and keeping it through year } j; k \in \{1, 2, \dots, K\}, i \in \{1, 2, \dots, N\}, \\ & j \in \{i, i+1, \dots, i+V-1\}
 \end{aligned}$$

For each individual objective $k = 1, 2, \dots, K$, we formulate the vehicle replacement problem as follows:

$$(2.1) \quad \text{Minimize} \quad \sum_{i=1}^N \sum_{j=i}^{V+i-1} c_{kij} x_{ij}$$

Subject to

$$(2.2) \quad \sum_{j=1}^{V-1} x_{1j} = 1$$

$$(2.3) \quad \sum_{j=i}^{V+i-1} x_{ij} \leq 1, \quad \forall i = 2, \dots, N$$

$$(2.4) \quad \sum_{j=i}^{V+i-1} x_{ij} = \sum_{n=1}^{i-1} x_{n,i-1}, \quad \forall i = 1, \dots, N$$

$$(2.5) \quad x_{ij} \in \{0, 1\}, \quad \forall i = 1, \dots, N; j = 1, \dots, (N+V-1)$$

The objective (2.1) is to minimize the total cost or environmental burden (specified by index k) of owning a vehicle for the time horizon. The first constraint (2.2) ensures that a vehicle is purchased in the first year of the horizon (no vehicle may be held over from a previous ownership). The next constraint (2.3) ensures that multiple vehicles may not be purchased in the same year. Finally, (2.4) ensures that a new vehicle is purchased as soon as one is retired, and (2.5) states that the variable x_{ij} is binary (0/1). This formulation of the model is solved using AMPL/CPLEX [50], and we use this formulation extensively in Section 3.3.2 to check the Pareto optimality of solutions found by the multiobjective metaheuristics.

2.3.2 Dynamic Program Formulation

An alternative formulation of the problem was implemented for our initial, numerical trials. For each objective that we desire to minimize (initially economic cost, energy use, *NMHC* emitted, *CO* emitted, *CO*₂ emitted and *NO*_x emitted), we construct a deterministic dynamic program to optimize over a finite horizon of predicted data. At the start of each year of ownership, the vehicle owner has the opportunity to **buy** a new vehicle or **keep** the existing vehicle. Each solution, therefore, is a series of replacement decisions of the form “Buy a new vehicle at the start of 2001 and keep it for 5 years; then buy a new vehicle at the start of 2006 and keep it for 8 years,” etc. We use the following notation in this dynamic programming formulation of the single-vehicle, single-objective model:

n = First year of the study

N = Last year of the study

M = Maximum physical life of a vehicle

$B_M(i)$ = Burden of the materials production of model year i vehicle

$B_A(i)$ = Burden of the manufacturing and assembly of model year i vehicle

$B_U(i, j)$ = Burden of the vehicle use during year j of model year i vehicle's service

$B_R(i, j)$ = Burden of the maintenance and repair during year j of model year i vehicle's service

$B_E(i, j)$ = Burden of the retirement of model year i vehicle at the end of year j

$u(i, j)$ = Burden of purchasing a new vehicle at the start of year i and keeping it for j years

x_i = Decision variable representing the number of years to own vehicle of model year i

For any i , $u(i, 0) = 0$ and represents the case in which a new vehicle is not purchased in year i . Therefore,

$$u(i, j) = \begin{cases} B_M(i) + B_A(i) + B_E(i, i + j - 1) + \sum_{k=1}^j (B_U(i, k) + B_R(i, k)) & \text{if } j > 0 \\ 0 & \text{if } j = 0 \end{cases}$$

For each criterion, this model seeks to minimize the burden from the life cycle of model years n to N by deciding how long to keep each vehicle before purchasing a new vehicle. For the model, vehicles owned for 0 years represent vehicle models that are never purchased. The dynamic programming optimality equations are constructed as follows. Let $f(i)$ be the minimum possible burden accumulated from the start of year i through the end of year N . Then

$$f(i) = \begin{cases} \min_{x_i \in \{1, 2, \dots, M\}} \{u(i, x_i) + f(i + x_i)\} & \forall i = n, \dots, N \\ 0 & \forall i > N \end{cases}$$

To solve each of these dynamic programs, the method of forward reaching [18] is implemented in a C program. For each program, in order to observe the behavior of the objectives not currently optimized, the values of these objectives under the generated policy are printed in addition to the optimal value of the single, minimized objective. From our initial runs of these programs, we are able to observe the expected tension between, for example, minimization of cost (few vehicle replacements) and NO_x emissions (several replacements in rapid succession to take advantage of technological improvement in fuel economy).

2.4 Single-objective Results

We began the Life Cycle Optimization of Vehicle Replacement study by considering economic cost, energy use, and emissions of nonmethane hydrocarbons ($NMHC$), nitrogen oxides (NO_x), carbon monoxide (CO) and carbon dioxide (CO_2). The optimal replacement policies for each of these are given in Table 2.1.

Table 2.1: Single-objective Optimal Policies (1985 - 2020)

Objective	Optimal Replacement Policy
Economic Cost	18, 18
CO_2	18, 18
Energy	18, 18
CO	3, 3, 4, 6, 6, 7, 7
$NMHC$	6, 6, 10, 14
NO_x	5, 5, 6, 6, 14
$NMHC + NO_x$	5, 5, 7, 7, 12

2.4.1 Refinement of Objectives

From these preliminary results we see that there is indeed tension among the objectives. We see that information regarding trade-offs between objectives would be interesting and helpful, but we first revisit our choices in objectives to streamline the future comparisons. We note that energy use and carbon dioxide (CO_2) emissions are nearly identical in profile since both are driven predominately by the combustion of gasoline. As changes in fuel economy will therefore affect these objectives simultaneously, we choose to eliminate the energy use objective, and represent this perspective with the CO_2 emission objective alone.

In our discussions with our research partners at General Motors, we discover that great improvements in the recent past have reduced the carbon monoxide (CO) emissions to nominal levels, and hence the introduction of new technologies would not be expected to impact this objective significantly. We therefore choose to eliminate this objective from further consideration as well.

Finally, we have discovered that it is a common practice to add the nonmethane hydrocarbon ($NMHC$) emissions and the nitrogen oxides (NO_x) emissions together, allowing the sum to sufficiently describe the impact of both pollutants since they work together in the formation of smog. This step is our final refinement to the portfolio of objectives considered, and it leaves us with three areas of analysis: discounted

economic cost, CO_2 emissions (a global pollutant), and $NMHC + NO_x$ emissions (local pollutants).

2.4.2 Objective Trade-offs

The single-vehicle results above in Table 2.1 show that the economic cost objective and the CO_2 emission objective are both minimized when a policy that mandates two 18-year vehicle lives is implemented. This contrasts with the $NMHC + NO_x$ policy of 5,5,7,7,12 years, and in Figure 2.1, we see the results of these policies on the nonminimized objectives. In this Figure, Policy 1 is the (5,5,7,7,12 years) that minimizes $NMHC + NO_x$, and Policy 2 is the (18,18 years). To move from Policy 1 to Policy 2 requires an increase in $NMHC + NO_x$ emissions of almost 50% over the 36 year study horizon, and the economic cost increase required to change from Policy 2 to Policy 1 exceeds 50% of the Policy 2 economic cost. The relative magnitude of the change on the CO_2 emissions is not great, and this is due to the fact that the CO_2 is primarily emitted in the combustion of gasoline, and this does not change much in the expected technology changes through 2020.




Figure 2.1: Magnitude of Trade-offs

CHAPTER III

Multiobjective Metaheuristics

From the single-objective analysis discussed in Chapter II, we discover that a single, optimal solution to our multiobjective, vehicle-replacement problem does not exist. If we wish to minimize the economic cost required to own and drive a vehicle, then we should choose to keep the vehicles in services for long periods of time (15 to 20 years). However, if we would like to minimize the nonmethane hydrocarbons and nitrogen oxides emitted, we should replace a vehicle more frequently so that new, efficient technology may replace older, possibly-failing, emissions-control systems. In order to present meaningful information regarding the impacts of various policies to a decision-maker, we look to the discipline of multiobjective decision-making.

Several methods for solving multiobjective dynamic programs exist in the literature, but the fact that our objectives are noncommensurable prevents easy implementation of these methods. We choose not to assign a dollar value to the various emissions since the true cost of these environmental burdens is not yet understood. Therefore we cannot simply sum the objectives and minimize one grand cost. In general, we seek Pareto-optimal solutions to the objectives such that changing the solution to improve in one area will cause decline in another. The sensibilities of the decision-maker will then help determine which of these nondominated solutions is

most appropriate.

One method of analyzing several objectives together involves the use of multi-attribute utility functions. Carnahan and Thurston [8] employ this method to determine the optimal set of design characteristics for a production process so as to optimize cost, pollution and quality. The construction of this utility function, however, requires that the decision-maker assign weights of importance to the objectives. For a single decision-maker this is perhaps manageable, but we seek to construct a model that produces a solution which is informative for our anticipated audience of consumers, automobile manufacturers and policymakers. Each of these would most likely weigh the objectives differently, and hence a solution without the preliminary weights may be more helpful.

Another method commonly used to evaluate problems with several objectives is the ϵ - Constraint Method. This method involves choosing one objective to minimize while constraining the others to desired values. By varying these constraint values, one is able to explore the possible solutions. This ϵ - Constraint technique is employed by Abo-Sinna and Hussein [3] as well as Wang and Shieh [64] who introduce a method of multiplier updating to an existing Iterative Dynamic Programming method for studying chemical processes. Ko, et al. [56] use this method in combination with a weighing approach to study trade-offs between cost, reliability and storage at a water reservoir. These techniques all suffer from the disadvantage of providing only one new nondominated solution at each iteration. For a complex problem in which the nature of the solution space is not known up front, these methods may require a large number of iterations before an adequate variety of solutions is found.

A third area of multiobjective solution techniques is evolutionary heuristic algorithms. These methods have the advantage that they maintain a population of

solutions throughout their implementation. This allows for the possibility of finding multiple solutions per iteration in addition to ending with a set of possible solutions for each run. For the present multiobjective, vehicle-replacement problem, we have not restricted our analysis to the wishes of a particular decision-maker. We prefer to construct a set of efficient solutions that demonstrate the various trade-offs required for each policy choice, and therefore the population-based, evolutionary, heuristic algorithms are an appropriate and attractive choice. In the following section we discuss some background regarding single-objective heuristics so that we may build upon these ideas in Section 3.2.

3.1 Single-objective Heuristic Algorithms

The two foundational areas for the algorithms to be discussed in Section 3.2 are genetic algorithms and tabu search. As the respective multiobjective algorithms are extensions of previous single-objective bodies of literature in both cases, we present some introductory ideas for each of these to provide a foundation for our later discussion.

3.1.1 Genetic Algorithms

The development of genetic algorithms and a number of other evolutionary programming techniques are described well in [31], [14] and [46], and we note that these are largely based on ideas given in [33]. These algorithms are most often used for large-scale or complex searches, and they are designed to mimic the processes of natural selection and genetics. A typical algorithm begins with a randomly generated *population* of potential solutions to the problem at hand. For ease of manipulation within the algorithm, each solution is *encoded* into a string format, according to a chosen scheme. The strings are referred to as *chromosomes*, after the genetics precedent,

and each component of a chromosome is called a *gene*, which may take any of the specified possible values known as *alleles*. For example, consider the problem

$$\text{Maximize } \{f(x) = x : x \in \{0, 1, 2, \dots, 31\}\}$$

[48]. The solutions to this problem may be encoded into 5-digit binary chromosomes in which each gene assumes the standard binary value. Hence the value 1 is encoded as 00001, the value 2 is encoded as 00010, and so on through the value 31 which is encoded as 11111. For each chromosome we calculate an associated *fitness* value according to the given objective function to be used in comparing solutions during the evolutionary process. Given the objective function above, for example, the fitness of individual 01100 would be 12.

Once the population is formed, a “survival-of-the-fittest” scheme is used to evolve this population into a new *generation*. As there is a direction in which the evolution is desired to progress (toward the optimal solution of the stated objective function), genetic operators designed specifically to guide the evolutionary process in this direction are applied. A typical algorithm will begin with an operator known as *reproduction* [31]. This involves the *selection* of *parents* from the current generation according to their fitness values so that the “better” genetic material will be used in construction of the next generation. The chosen parents are copied into a *mating pool* which will undergo further genetic operators. A number of alternatives for this selection process have been evaluated, and again we refer the reader to [31], [14] and [46] for additional reference.

Once this mating pool is constructed, a *crossover* operator is typically applied to randomly selected pairs of parents within the pool. For *single-point* crossover, once two strings are chosen, a point on the chromosome is randomly chosen such that the genes past this point are swapped on the two strings. For the example above, if we

begin with two parents

10111 and 01100,

and if we happen to choose to crossover from the third gene on, the resulting new solutions are

10100 and 01111.

As in the case of the reproduction operator, there have been several successful variations on the single-point crossover [46], including extensions to *2-point* and *multi-point* crossover. Reproduction and crossover operators constitute the majority of the evolutionary process for many genetic algorithms [31], but a final kind of operator is employed to ensure that pieces of the genetic picture (the possibility of a particular allele in a given position on the chromosome, for example) are not lost. In most cases this final operator is *mutation*. A few solutions in the mating pool are randomly selected to have single genes randomly replaced with small probability by alternative allele values, and this helps maintain diversity within the population. Once the new generation is complete, this becomes the next population from which to select parents, and the process is repeated for a specified number of generations or until a sufficiently “good” fitness value is found.

Single-objective genetic algorithms have been applied with success to a variety of problem classes. Michalewicz [46], for example, discusses genetic algorithm implementations for the Transportation Problem, the Traveling Salesman Problem, and a number of graph-drawing, scheduling and partitioning problems. The theory of convergence for single-objective genetic algorithms has also been studied extensively, see for example [54]. The extension of genetic algorithms to problems which consider multiple objectives is a more recent area of research, and in Section 3.2 we will discuss the development and use of genetic algorithms for our multiobjective vehicle

replacement problem. In Section 3.1.2 below, we introduce background material on tabu search algorithms.

3.1.2 Tabu Search

Tabu search is a metaheuristic for solving combinatorial optimization problems that employs flexible and dynamic memory structures to aggressively explore past the limitations of local optimality ([30], [28]). For a thorough study of tabu search, see [24]. For the purposes of this discussion, we consider the following problem

$$\text{Minimize } f(x) : x \in X \text{ in } R_n$$

where $x \in X$ constrains x to a discrete set of values. We define the set $S(x)$ to be the *moves* $s \in S$ available to a solution x , where $s : X(s) \rightarrow X$. That is, $S(x) = \{s \in S : x \in X(s)\}$, and we refer to the set $S(x)$ as the *neighborhood function* [28].

For a typical implementation of tabu search, we begin with an initial $x \in X$, and we define a subset T of S so that

$$T(x) = \{s \in S : s(x) \text{ violates the given tabu conditions}\}.$$

These *tabu conditions* are a restriction on the moves to new solutions that a solution, x , may take in its future iterations. The purpose of the tabu list is to use historical information in the search process to prevent the algorithm's returning too soon to previously explored solutions. The algorithm begins with T empty, and from x it searches for the $s \in S(x) - T$ with the "best" possible characteristics, as defined by a pre-specified evaluation function, and we let this be the move that takes x to \bar{x} . We replace x by \bar{x} , and if $c(\bar{x}) < c(\text{the best solution found so far})$, then we update this best-found indicator. We then update the tabu list, T , according to its definition,

and we repeat this process until the chosen number of iterations are completed or until $S(x) - T = \emptyset$ [28].

Just as advances and alternatives have been developed for the genetic algorithm operators discussed above, the schemes for defining tabu lists and the “best” neighbor functions have been refined and adapted for a variety of problem-specific advantages [24]. Glover [29] describes tabu search applications in graph theory problems, course scheduling, telecommunications, flow shop sequencing, the Traveling Salesman problem, quadratic assignment, and character recognition. In the sections below we discuss the extension of tabu search as well as genetic algorithms our multiobjective, vehicle-replacement problem.

3.2 Multiobjective Heuristic Algorithms

Beginning with the publication of Schaffer’s Vector Evaluated Genetic Algorithm (VEGA) in 1984 [53], there have been several approaches to the creation of a set of Pareto-optimal solutions by means of a genetic algorithm. The success of early algorithms such as the Niche Pareto Genetic Algorithm (NPGA) [35], the multiobjective genetic algorithm introduced by Fonseca and Fleming [26] and the original Nondominated Sorting Genetic Algorithm (NSGA) [57] demonstrated the effectiveness of such population-based heuristics for finding a representative set of Pareto-optimal solutions.

As these and many other multiobjective algorithms have been studied and implemented, a new generation of algorithms has emerged. Zitzler and Thiele [66] introduce the Strength Pareto Evolutionary Algorithm (SPEA) which, in addition to their new niching method, powerfully combines several previously isolated ideas: external storage of nondominated solutions, assignment of fitness based on Pareto

dominance, and clustering to reduce the number of nondominated solutions. Given the strong performance by SPEA on several test problems ([41], [22], [66]), we choose to implement a variation of it, SPEA/M, for our multiobjective equipment replacement problem, in which the addition of the ‘M’ is used to denote its development at the University of Michigan.

Another multiobjective genetic algorithm that builds on earlier algorithms is NSGA-II [40], the new version of the Nondominated Sorting Genetic Algorithm. The elimination of the previously sensitive sharing parameter and the modification to an elitist selection process have addressed many of the issues regarding the original NSGA. Due to the strength of these updates, we choose to implement a variant of NSGA-II, the NSGA-II/M, as the second multiobjective genetic algorithm for the vehicle replacement problem considering environmental and economic objectives. For a more comprehensive survey of multiobjective evolutionary algorithms and their development, please see [61], [12] or [27].

Although, as described in Jones et al. [19], 70% of the multiobjective meta-heuristic literature employs genetic algorithms, there is also strong representation by simulated annealing and tabu search algorithms. We choose to expand our analysis of multiobjective solution techniques by considering a tabu search-based approach in addition to the evolutionary algorithms introduced above. Given the performance of Hansen’s Multiobjective Tabu Search (MOTS) [32] for project scheduling as described by Viana and de Sousa [62], we choose to develop MOTS/M to introduce the ideas from MOTS to the multiobjective vehicle replacement problem. As we introduce the details of the three algorithms presented in this paper, it is worth mentioning our intended framework for evaluating the effectiveness of the algorithms. Deb [17] states “it can be conjectured that there are two goals in a multiobjective optimiza-

tion: 1. To find a set of solutions as close as possible to the Pareto-optimal front [and] 2. To find a set of solutions as diverse as possible.” This two-fold goal of convergence and spread is our guide as we implement the following heuristic algorithms for the Life Cycle Optimization of Vehicle Replacement problem.

3.2.1 Strength Pareto Evolutionary Algorithm (SPEA/M)

Our variant of Zitzler and Thiele’s SPEA [66] for the Life Cycle Optimization of Vehicle Replacement begins with a randomly generated population, P , of predetermined size N . Any infeasible solution created at any time by the algorithm is discarded and replaced by another randomly generated solution. The locally noninferior members of P are copied and appended to an external set P' (initially empty). We then remove any duplicates or newly inferior solutions from P' . If the size of the set P' exceeds its predetermined limit, N' , we use the clustering procedure suggested by Zitzler and Thiele [66] to reduce the size of P' to N' .

The next step is to calculate, for each solution, a fitness value to be used in the creation of the next generation. The fitness values are constructed to encourage both the favoring of locally noninferior solutions and the spreading out of solutions along the Pareto frontier. For each solution $x \in P'$, let $n(x)$ be the number of solutions in P which are inferior to x . We assign to each $x \in P'$ a *strength*, $s(x) = \frac{n(x)}{N+1}$, and we use this strength as the measure of fitness for the solution. A small strength value indicates that the solution is isolated, and we favor these solutions by minimizing strength in the operations performed on P' . For each solution $y \in P$, let P'_y be the set of $x \in P'$ to which y is inferior. We assign to y a fitness $f(y) = 1 + \sum_{p \in P'_y} s(p)$. Again, the isolated solutions will be favored by minimizing fitness values.

Once the fitness values are assigned, we construct a mating pool of size N from which to create the next general population. We set aside a predetermined number,

Table 3.1: SPEA/M Pseudocode

1	Initialize P, P'
2	Find noninf. solns in P
3	Append to P'
4	Assign fitness to P, P'
5	Select N from $P \cup P'$
6	Crossover, Immigration
7	Repeat from 2

$\nu < N$, of places in the mating pool to be filled later, and we begin by using binary tournament selection with replacement on the union of P and P' to fill the first $N - \nu$ places in the mating pool. Binary tournament selection is an operator that randomly chooses two solutions from $P \cup P'$, compares their fitness values, and keeps only the better of the two. After $N - \nu$ comparisons, we have a pool of $N - \nu$ solutions to which we apply Bernoulli (also called uniform) crossover. Finally, the remaining ν places in the next generation are filled by the *immigration* [6] of ν new, randomly-generated solutions as an alternative method to mutation for the introduction of new genetic material to the population. The new generation, P , is now created. Unless we have completed the specified maximum number of generations, we begin again by appending copies of the locally noninferior solutions found in P to P' . Once all generations are completed, the output to the decision-maker is P' , the external list of noninferior solutions that has been updated and maintained throughout the course of the algorithm. For an outline of the SPEA/M, see Table 3.1.

3.2.2 Nondominated Sorting Genetic Algorithm II (NSGA-II/M)

The implementation of our variant of NSGA-II described by Deb et al. (2002) also begins with a randomly generated population, P , of size N . As described above, we discard and replace any infeasible solutions which are created at any time. In this case, however, there is not a separate list of locally noninferior solutions. We begin by finding the locally noninferior solutions within P and assigning these a rank

of 1. They are the first noninferior front found within P , and for a minimization problem, these would be the lightly-shaded points in Figure 3.1. We then search P again, leaving out the solutions in the first front. The set of locally noninferior solutions found in the remaining subset of P (the unshaded points in Figure 3.1) are assigned a rank of 2, and this second front is then set aside while we again search the remaining subset of P for a third locally noninferior front (in Figure 3.1 this is only the darkly-shaded point). This process continues until all solutions in P have been assigned a rank. The next step is to create a new population, Q , of size N from P .




Figure 3.1: Pareto-optimal fronts sorting

We apply to P the same binary tournament selection with replacement, Bernoulli crossover and immigration operators that were used in SPEA/M. In the first generation, the noninferior front ranks assigned to the solutions in P are used for comparing individuals in the tournament selection. In later generations, the evaluation criterion is the crowded-comparison operator described below. Once Q is created, we form a new population R , of size $2N$, by simply combining the populations P and Q .

For each member of R , we now calculate a crowded-comparison value using the

following three steps. First, we sort R into its locally noninferior fronts and assign ranks as described for P above. Within each front, Γ , we then calculate the *crowding distance* [40] for each solution to give a measure of the density of solutions surrounding it. For each $x \in \Gamma$, we calculate $\Gamma[x]_{distance(k)}$ for each objective k , and the sum of these distances is the crowding distance for x . Let $n = |\Gamma|$, and let $\Gamma[x]_k$ be the value of objective k for solution x . To compute the $\Gamma[x]_{distance(k)}$ values:

- For each $x \in \Gamma$, set $\Gamma[x]_{distance(k)} = 0$
- Sort Γ according to objective k so that $\Gamma[1]_k$ is the least objective value, and $\Gamma[n]_k$ is the greatest
- Set $\Gamma[1]_{distance(k)} = \Gamma[n]_{distance(k)} = \infty$
- For $i = 2$ to $(n - 1)$, $\Gamma[i]_{distance(k)} = (\Gamma[i + 1]_k - \Gamma[i - 1]_k) / (\Gamma[n]_k - \Gamma[1]_k)$

Once crowding distances have been assigned within all locally noninferior fronts, we construct a final ordering of R . We sort R first in increasing order of the locally noninferior front rankings, and then within each front, we sort according to the crowding distances just calculated. Because we wish to favor the isolated solutions, we prefer the solutions with the largest crowding distances. The last step in the algorithm is to create a new population, P , with the N best members of R . We then either repeat, beginning with the generation of Q , or quit if the desired number of generations is completed. A summary of the steps is given in Table 3.2.

3.2.3 Multiobjective Tabu Search (MOTS/M)

The Multiobjective Tabu Search introduced in [32] simultaneously optimizes a set of individual current solutions in order to produce a set of noninferior solutions, similar to the output of the multiobjective evolutionary algorithms. For MOTS/M and the

Table 3.2: NSGA-II/M Pseudocode

1	Initialize P
2	Sort to noninf. fronts
3	Assign ranks
4	Select Q from P
5	Form $R = P + Q$
6	Calc crowding distances
7	Sort R
8	Choose best N as new P
9	Repeat from 2

Life Cycle Optimization of Vehicle Replacement problem, we begin with a randomly generated set, X , of solutions (infeasible solutions are discarded and replaced) which we address one at a time in each iteration of the algorithm. The algorithm employs a weight vector, λ , initialized to 0 each time a new solution is considered, and a vector of range equalization factors, π , used to ensure valid comparisons across the K objectives. The range equalization factors are initialized to $\pi^k = 1/K$, for each objective $k = 1, \dots, K$, at the start of each solution's iteration. The MOTS/M, in the same way as the SPEA/M, maintains an external set of the locally noninferior solutions found so far. This set is initially empty, and it is augmented throughout the algorithm rather than overwritten. Additionally, let $f(x) = (f_1(x), \dots, f_K(x))$ be the vector of objective functions to be evaluated, all of which we assume are to be minimized.

For each solution $x \in X$, in each iteration, we begin by calculating the parameters relating it to the other solutions in X . Iteratively, for each solution $y \in X$ such that $f(y)$ is nondominated by $f(x)$ and $f(y) \neq f(x)$,

1. Set $w = 1/(\sum_{k=1, \dots, K} \pi^k |f_k(x) - f_k(y)|)$.
2. For each objective k , if $f_k(x) < f_k(y)$, then set $\lambda^k = \lambda^k + \pi^k w$.

If the resulting λ is 0, then we set each $\lambda^k = 1/K$; otherwise, we normalize the λ found above. The next step is to find the “best” neighbor of x to replace it in the

set of solutions. For each solution $x \in X$, we maintain a list of neighbors from which to choose. As described more fully in Section 3.3 below, the solutions for the Life Cycle Optimization of Vehicle Replacement problem are binary strings of Buy/Keep replacement decisions, and a solution's neighborhood is the set of solutions which can be attained by a single swapping of two consecutive decisions. As an example, if a particular solution, x , has "Buy" in period 4 and "Keep" in period 5, one of its neighbors is the solution whose only difference from x is that it requires a "Keep" in period 4 with the "Buy" in period 5; that is, we delayed the replacement by one period.

To find the best neighbor, we consider two criteria. First, the move from x to the neighbor must not be on the tabu list for x . This tabu list is a set of the most recent moves (from the corresponding solution to a neighbor) that we want to keep from repeating continuously. The length of the tabu list (the number of forbidden moves) is a parameter of the algorithm, and when the most recent move is added to the list, the earliest move is removed. The second criterion is to find the neighbor, y , that minimizes $\lambda \cdot f(y)$. Once the best neighbor, y^* , is found, we add the move (y^* -to- x) to the tabu list, removing the earliest move from the list if necessary. We then replace x by y^* in the set of solutions and check to see whether this new x is noninferior with respect to the externally kept set of locally noninferior solutions. If x is noninferior, we add it to this external set, and update the set to remove any newly inferior solutions. If the noninferior set contains more than two solutions, we update the range equalization factors,

$$\pi^k = \frac{1}{\text{Range}^k} \left[\sum_{i=1}^K \frac{1}{\text{Range}^i} \right]^{-1}$$

where Range^k is the range of values of objective k within the noninferior set, and K is the number of objectives.

Table 3.3: MOTS/M Pseudocode

1	Initialize set X
2	For each $x \in X$:
3	Update w and λ
4	Choose best neighbor y
5	Update tabu list
6	Test inferiority of y
7	Update π
8	Inc. drift counter
9	Repeat from 2

Finally, we increment a *drift* counter which, when it has reached the parameter-specified threshold, indicates the interval at which we randomly select one solution in the general set to be replaced by another randomly chosen solution in the set. We then move to the next solution in the set and repeat the procedure until the specified number of iterations are completed. An outline of the MOTS/M algorithm is given in Table 3.3.

3.3 Implementation

With the aid of multiobjective metaheuristics, we are able to construct a set of Pareto-optimal solutions to the Life Cycle Optimization of Vehicle Replacement problem discussed in Chapter II. Each solution represents a replacement policy for a single vehicle over the given time horizon. Analysis of this Pareto-optimal set provides information on the trade-offs among the objectives and allows a decision-maker to choose a final policy according to his specific agenda.

The three objectives we wish to minimize are economic cost (constant 1985 dollars), carbon dioxide (CO_2) emissions, and the sum of nonmethane hydrocarbons and nitrogen oxides ($NMHC + NO_x$) emissions. Preliminary individual analysis of these objectives demonstrates that the $NMHC + NO_x$ (the *local* pollutant objective) changes almost inversely proportionally to the economic cost and CO_2 (the *global*

pollutant objective). From the results of Chapter II we have seen that as technology improves, newer-model vehicles are expected to emit fewer local pollutants, so to minimize these emissions one should replace a vehicle fairly frequently, taking advantage of the fact that with each replacement the local emission performance will improve. On the other hand, since carbon dioxide emissions and economic operating costs are closely related to annual fuel consumption, the recommended replacements for these objectives are farther apart due to the more modest predictions for widespread fuel economy improvement. This tension among the objectives indicates that further information regarding trade-off among the objectives may be gleaned from the construction of a set of Pareto-optimal solutions for the simultaneous minimization of all three objectives.

3.3.1 Novel Algorithm Details

The algorithm outlines given above in Sections 3.2.1 - 3.2.3 explain our adaptations of three published multiobjective heuristic algorithms for the vehicle replacement problem ([66], [40], [32]). In all cases we used binary encoding for the replacement policies so that for our time horizon of 36 years (1985 - 2020), each solution is a binary string of length 36. We use “1” to indicate that a purchase occurred in the corresponding year, and a “0” to indicate that the existing vehicle was kept. All policy strings begin with a “1”, as we require a vehicle to be purchased at the start of the time horizon (1985). The only additional feasibility constraint is that no vehicle may be kept for more than 20 years. This is implemented by discarding any solutions that violate this maximum life constraint before they are added to any of the solution populations.

For the SPEA/M and NSGA-II/M, we introduce two genetic operators not used with either of the two original algorithms. Traditionally, mutation has been used

to introduce new genetic material to a population [31]. This operator chooses a generally small set of individuals each generation to undergo random alterations to their encoded representations. As an alternative to mutation, we choose to employ *immigration* [6] to introduce new possibilities in each generation. Instead of altering existing solutions, the immigration operator randomly generates completely new solutions to be added to a specified number of places in the generation’s population. The second modification we make to the original genetic operators for SPEA/M and NSGA-II/M is to consistently use Bernoulli (also called *uniform*) crossover [31] in the evolution of each generation. This means that for each crossover operation we select two individuals and randomly decide “switch” or “don’t switch” for each chromosome (rather than choosing just one place on the string to trade genetic information).

Finally, as discussed in the following section, we implement a method to validate our final set of solutions from each algorithm. This is not addressed in the papers which introduce our foundational algorithms ([66], [40], [32]), as these used known functions for their test problems. In this case, we do not know the solution set that we seek as we begin our analysis of the multiobjective vehicle problem, so we must validate the results after the fact.

3.3.2 Verifying Pareto Optimality

One drawback of the multiobjective heuristic algorithms is that the set of locally noninferior solutions given as output might contain solutions that are not truly Pareto-optimal. For the Life Cycle Optimization of Vehicle Replacement problem, we construct an integer program to verify the Pareto optimality of each solution given by the algorithms. This method may be used to verify Pareto optimality of any heuristic, hence we consider it separately rather than as part of our heuristic algorithms.

Let X be the set of feasible solutions to the problem, K the number of objectives to be minimized, and T the number of years in the time horizon. For each output solution, $x = (x_1, x_2, \dots, x_T)$, with objective vector $f(x) = (f_1(x), \dots, f_K(x))$, we solve the following:

$$(3.1) \quad \text{Maximize} \quad z = \sum_{i=1}^K (f_i(x) - f_i(y))$$

Subject to

$$(3.2) \quad f_i(y) \leq f_i(x) \quad i = 1, \dots, K$$

$$(3.3) \quad y \in X$$

If $z > 0$, then we have found a solution, y , that was able to make an improvement over x in at least one objective without degrading the values of any other objectives. That is, $z > 0$ implies that x is not Pareto-optimal, and $z = 0$ indicates that we were not able to improve on x which is therefore proven to be Pareto-optimal. If we find that a particular solution $\bar{x} \in X$ is not Pareto-optimal, we can then use the constraint method [10] to try to find Pareto-optimal solutions in addition to z . Considering each objective component $k \in [1, 2, \dots, K]$ in turn, we solve

$$(3.4) \quad \text{Minimize} \quad z_k = f_k(x)$$

Subject to

$$(3.5) \quad f_i(x) \leq f_i(\bar{x}) \quad i = 1, \dots, K$$

$$(3.6) \quad x \in X.$$

Because \bar{x} is not Pareto-optimal, we are guaranteed to find at least one solution x such that $f_k(x) < f_k(\bar{x})$ for some $k \in [1, 2, \dots, K]$ (though it could be only the y from above), and hence we will find at least one Pareto-optimal solution that may be added to our solution set.

Table 3.4: Pareto-optimal solutions found in 100 generations of the SPEA/M

	Minimum	Maximum	Median
500/50/100/2	31	48	36.5
500/50/100/5	30	47	37.5
500/50/400/2	45	49	47
500/50/400/5	40	49	46.5
500/250/100/2	31	49	38
500/250/100/5	30	51	38
500/250/400/2	57	69	61.5
500/250/400/5	53	67	61.5

The validation of every solution found is a significant task. However, the implementation of the heuristic optimization followed by the integer program validation of Pareto optimality does guarantee the quality of the solutions presented to the decision-maker. In the experiments discussed below, we verified the Pareto optimality of the solutions presented, but we did not extend to finding new solutions with the constraint method.

3.3.3 Algorithm Performance

SPEA/M Results

We tested several parameters for the SPEA/M, including the initial population size, the limit on the size of the external set of noninferior solutions, the number of generations, and the numbers of crossover and immigration to perform in each generation. For each parameter set we conducted 10 trials, each beginning with a different random seed. Hence median results are presented for the output. To represent the parameters used in each trial, we adopt the following notation: N/N' /Number of crossover/Number of immigration, where N is the general population size, and N' is the maximum size of the external set of noninferior solutions. An example of preliminary work to evaluate the parameter choices is given in Table 3.4. As we are able to obtain such a large set of Pareto-optimal solutions, we are able to see a visual representation of the trade-offs among the objectives by plotting the objective values

in pairs. For one of the instantiations of the 500/250/400/2 parameter set after 100 generations, see Figure 3.2.

We observe a confirmation of the tension between the $NMHC + NO_x$ objective and the other two that had been identified in initial single-objective analysis. An interesting feature of the Pareto-optimal front is the clustering of the solutions into four regions, particularly evident in the $NMHC + NO_x$ vs. Cost plot. Each cluster corresponds to a particular number of vehicle replacements over the time horizon. For example, the solutions with costs less than \$50,000 each require an initial purchase and just one replacement over the horizon. The next cluster, containing solutions with costs between \$50,000 and \$60,000, represents Pareto-optimal replacement policies which require two replacements after the initial purchase. The third and fourth clusters correspond to policy groups requiring three and four replacements after the initial purchase, respectively.

Additional analysis using the Constraint Method has verified that there do exist these gaps in the cost objective within which no Pareto-optimal solutions can be constructed. The threshold areas immediately before and after the jumps in the costs provide interesting information to the decision-maker. In most cases, a cost occurring just above a gap could be reduced significantly without great impact on the other two objectives. Although the large number of Pareto-optimal solutions found is most likely more information than a decision-maker would need, the insight into the objective trade-offs gained by viewing so much of the Pareto-optimal front could enhance the understanding of such a decision-maker.

NSGA-II/M Results

We now examine the results from the NSGA-II/M for the Life Cycle Optimization of Vehicle Replacement problem. Similar to our approach with the SPEA/M, we tested






Figure 3.2: Pareto-optimal solutions after 100 SPEA/M generations

Table 3.5: Pareto-optimal solutions found in 100 generations of the NSGA-II/M

	Minimum	Maximum	Median
500/100/2	63	67	65
500/100/5	64	69	66
500/400/2	65	68	67
500/400/5	64	68	67

several parameter options, the output for each set being a representation of the 10 random trials performed. As there is no external set of noninferior solutions for the NSGA-II/M, we have one fewer parameter than for the SPEA/M, and we use the notation (Population size/Number of crossover/Number of immigration) to specify a set of trials. An example of the output from the NSGA-II/M for parameter choices similar to those used for the SPEA/M is given in Table 3.5. This performance appears comparable to the SPEA/M, and we explore the similarities further in Section 3.3.4. Again we would like to visually observe the relationships between the objectives as demonstrated by the Pareto-optimal solutions found by the NSGA-II/M. We choose one set of output for the 500/400/2 set of parameters to plot for demonstration in Figure 3.3. We observe that the NSGA-II/M has identified the four clusters in the cost objective, and the decision-maker again has many solutions from which to choose a replacement policy.

MOTS/M Results

Our variant of the Multiobjective Tabu Search algorithm was not as successful as the two evolutionary algorithms. In addition to the fact that considering each solution individually led to slower progress, we were unable to converge to more than a few solutions on the true Pareto-optimal front. To illustrate this result, we plot in Figure 3.4 the output from one MOTS trial with a set of Pareto-optimal solutions found by the NSGA-II/M. In this MOTS/M trial we allow 500 initial solutions, 500 “generations,” a tabu list length of 2 moves for each solution, and a drift criterion of 1000

Table 3.6: Outline of the three algorithms

Stage	SPEA/M	NSGA-II/M	MOTS/M
Initialization	Population P of size N , empty P'	Population P of size N	Set X with neighborhoods and empty external set
Method to encourage diversity of solutions	Assignment of strength and fitness values	Crowding distance calculations	Multiplicative weights (λ)
Chance to select favorable solutions	Binary tournament selection on $P \cup P'$ to produce mating pool	Binary tournament selection on P to produce Q ; ranking of R	Choice of best neighbor
Recombination of solution fragments	Crossover	Crossover	Neighborhoods
Introduction of new solution fragments	Immigration	Immigration	
Advance the algorithm	Begin again with new P	Form new P from best N solutions in $R = P + Q$	Replace each solution one-at-a-time with best neighbor

iterations before drift is performed. Only two of the plotted MOTS/M solutions are Pareto-optimal, though several are “close” to Pareto-optimal solutions found by the NSGA-II/M. We are able to see two new clusters in the cost objective corresponding to policies that require 5 and 6 replacements after the initial purchase, but these are never policies whose objective values are nondominated.

3.3.4 Algorithm Comparison

Although the implementation details are varied, there exist functional commonalities between the three algorithms. We outline some of these in Table 3.6 to illustrate the similar behaviors. Because the SPEA/M and the NSGA-II/M both appeared to successfully find representative sets of the Pareto-optimal front, we choose to explore their performance in more detail. From the initial parameter-testing phases, we determine that an initial population of 500 solutions with 400 crossover and 2 immigration per generation is a favorable parameter set for both algorithms, and we let the maximum number of noninferior solutions in the SPEA/M be 250. We first examine the median number of Pareto-optimal solutions as the number of generations increases. A set of reference points of this progress are plotted in Figure 3.5. We

Table 3.7: Variance within the trials as generations increase

Generations	SPEA/M Variance	NSGA-II/M Variance
50	10.44	0.99
100	16.54	1.12
200	16.27	1.43
300	15.38	1.29
400	15.38	1.29
500	15.38	1.34
750	14.9	1.34

discover that although both algorithms perform well as the generations increase, the NSGA-II/M attains a greater median number of Pareto-optimal solutions found after just 50 generations than the SPEA/M finds after 750 generations. Additionally, as we examine the variance within each sample of 10 trials per experiment, we discover that the NSGA-II/M reaches its exceptional performance with a minimum of variance, as shown in Table 3.7.

3.4 Remarks and Conclusions

The Life Cycle Optimization of Vehicle Replacement problem is a complex problem with more than the commonly-considered two objectives. We have shown that the frameworks of the Strength Pareto Evolutionary Algorithm [66] and the Non-dominated Sorting Genetic Algorithm II [40] successfully identified large numbers of Pareto-optimal solutions. These solutions were sufficiently well-distributed along the Pareto-optimal frontier that we were able to distinguish clusters of solutions corresponding to characteristics of the replacement policies.

Although our variant of the Multiobjective Tabu Search [32] was not able to discern many truly Pareto-optimal solutions, we did approach a portion of the Pareto-optimal front. The modification of traditional tabu search to include a set of solutions for multiobjective optimization allowed us to introduce a nontraditional competitor for the multiobjective evolutionary algorithms.






Figure 3.3: Pareto-optimal solutions after 100 NSGA-II/M generations




Figure 3.4: MOTS/M solutions with Pareto-optimal solutions for comparison




Figure 3.5: Pareto-optimal solutions as generations increase

CHAPTER IV

Multivehicle Study

In the previous two chapters, we have examined extensively the service life of a single passenger vehicle. Using historical data and projections for the future, we have been able to observe tradeoffs among the environmental and economic objectives as we lengthen or shorten the vehicle's service life. Replacement decisions are not made in isolation, however, and we now consider passenger vehicle replacement within the nationwide vehicle population. Within this population we have households owning vehicles of various ages, and the replacement decisions in these households made be the result of a variety of motivations.

Vehicle type choice models have shown that factors such as household income and size, in addition to vehicle price, capacity, performance characteristics, age, brand, style, and safety features may affect a consumer's vehicle choice [11]. Given an existing population of households and their vehicles, however, how might the profile of this population change due to an alteration in its environment? We construct a model to examine the impact on a population of a change in national policy such as fuel cost or the widespread introduction of advanced technology.

4.1 Policy Analysis of Complex Systems

Public policy analysis is defined by Walker in [63] as “a rational, systematic approach to making policy choices in the public sector...[:] a process that generates information on the consequences that would follow the adoption of various policies.” He continues to describe the benefits of a sound relationship between an analyst working with the data and modeling aspects of a policy decision and the policymaker whose judgement helps define objectives and constraints [63]. As we consider policies that could impact the national population of vehicles, we see that describing the consequences requires gleaning information from an elaborate and interdependent system of stakeholders. These include consumers, automotive industry partners in manufacturing and R&D, energy industry leaders, and numerous public policymakers. In addition to considering multiple vehicles, we must also consider the effect of multiple owners in areas with diverse geographic profiles and a continually changing economic climate, as just a beginning. Hence we see the effect of deliberate *policy changes* as well as changes in the environment brought about by *external forces* [63].

Bankes ([4], [5]) highlights many problems that may result from the application of classical optimization techniques to analysis of policy in a complex environment. He notes, “For those problems for which no model can accurately predict the details of system behavior, approaches to policy analysis based on using some model to forecast system behavior will be inappropriate.” [5] He suggests alternatives to a single forecasting model which include consideration of an *ensemble* of plausible models to determine robust policies or level sets of solutions which perform well relative to a determined threshold [5]. This approach has been used, for example, to study the effect of greenhouse-gas reduction strategies on climate change [51].

The aim of using computer models, and well-reasoned ensembles of computer models in particular, to aid policy decisions is to gain insight into the possible future states of the system [4]. We began our study of the multiobjective, vehicle-replacement problem by developing of a set of Pareto-optimal solutions for single-vehicle replacement, and this allowed us to observe trade-offs among environmental and economic objectives in detail. We now construct a collection of scenarios that allows us to examine the impact of a changing national environment on a population of vehicles and consumers.

4.2 Model

To begin to understand the dynamics of the national vehicle-owning population, we work to group households into categories according to their perceived purchasing behavior. As the single-vehicle analyses gave insights specifically for 5-passenger, domestic sedans, we begin with this part of the vehicle population for our multivehicle study. We assume that many of the consumers driving 5-passenger, Primary Use sedans are interested in minimizing the cost to drive their vehicles for commuting, etc. With our focus on one kind of vehicle, we clearly can only attempt to understand the effect of a policy change on the owners of vehicles of this type. If we could characterize all existing households, we would be able to realize the model shown in Figure 4.1.

For the model illustrated in Figure 4.1, we begin with a population of vehicles and owners in equilibrium. Let there be N categories of households so that each household category $i \in 1 \dots N$ constitutes the fraction ω_i of the population. When the indicated Policy Change is imposed, each category of households may respond uniquely to this change with a behavior change denoted $\Delta_i, i \in 1 \dots N$. We weight

these changes according to the representation of each household in the entire pool (ω_i), and we aggregate the behavior changes to realize the overall impact of the policy change. While we could never presume to correctly determine the change for a specific individual, the goal is to find the appropriate derivative in the aggregate case.




Figure 4.1: Policy Decision Support System Framework

4.3 Data

Many of the assumptions made in Section 2.1 for the single-vehicle replacement model hold in the population-wide model as well. We continue to consider only Primary Use Vehicles, and within this subset of the population, only the midsize, 5-passenger, domestic sedans. Each vehicle is driven 12,000 miles per year, maintained according to manufacturer specifications, filled with regular unleaded gasoline, and retired within 20 years of its initial purchase.

Although the baseline vehicle used in this multivehicle analysis has the same profile as the vehicle discussed in Chapter II, we have an additional need for data

describing consumer behavior and alternatives resulting from policy changes.

4.3.1 Household Profile

We continue to analyze midsize, domestic, passenger vehicles driven 12,000 miles per year, and we begin our analysis assuming that such an owner is only concerned with the economic cost of his or her vehicle use. Our network of these households is described as follows:

- One midsize, 5-passenger, domestic sedan per household is driven 12,000 miles per year
- The consumer minimizes the cost-to-go as each year's buy/keep choice is made
- Ownership begins with a uniformly assigned 1-20 year-old vehicle at the end of 1994 (hence the model years may begin in 1975)
- Annual buy/keep decisions continue through 2050 to reduce end-of-study effects through 2020

4.3.2 Baseline Consumer Behavior

As introduced in Section 4.2, we assume that our population of vehicles and owners begins in some equilibrium state. That is, vehicles of the profile used in Chapter II are being driven and replaced by consumers according to the household profile discussed above. We refer to this instantiation of the model (with no policy interference) as the **Baseline** scenario, and the resulting household replacement schedules are shown in Table 4.1. The first column contains the age of the vehicle owned at the end of year 1994. These ages are randomly assigned to consumers, and over the time horizon 1995-2020, we see that the replacement policies vary according to this starting point. The second column displays the ages of vehicles owned by a baseline consumer who

begins with the corresponding vehicle age at the end of 1994. For example, in the first row we see that a consumer whose vehicle was one year old at the end of 1994 will keep that initial vehicle until it has been driven for 19 years. The vehicle will then be replaced by a vehicle which becomes eight years old in 2020, where the parentheses around the vehicle age at the end of year 2020 indicate that this would not have been the last year of the vehicle's life, had 2020 not been the final year of our study horizon. The third column shows the years in which vehicle replacements were made; all replacements occur at the beginning of the indicated year.

We see that the optimal behavior for all consumers was to keep the vehicles for 19 or 20 years before replacing with a new vehicle. Over the 1995-2020 period, the population-wide average vehicle age per year varied from 9.09 to 11.38 years, with an average of 10.02 years. The average age of automobiles in operation in the United States in 1997 was computed to be 8.7 years [15], and this would include effects of vehicles taken off the road following accidents as well as consumers with objective functions other than minimizing cost.

As we currently assume the consumers are minimizing their cost to own and drive a vehicle, we examine the average cost per year. For a population of 200 households, the average expenditure per year varies from \$1644.14 to \$7814.76, with the average annual expenditure over 1995-2020 of \$3683.20 (all costs given here in year 2000 dollars). The U.S. Bureau of Labor Statistics' Consumer Expenditure Survey estimates that households in 2000 spent an average of \$6990 on personal vehicle transportation [1]. However, in single-person households which are reported to have 1.0 vehicles per household on average, the average expenditure on personal vehicle transportation in 2000 was \$3410 [1], just below the \$3683.20 found in our baseline multi-household scenario.

Table 4.1: Baseline Consumer Behavior

Age at end of 1994	Ages of vehicles owned 1995-2020	Replacement Years
1	2-19, 1-(8)	1994, 2013
2	3-20, 1-(8)	2013
3	4-20, 1-(9)	2012
4	5-20, 1-(10)	2011
5	6-19, 1-(12)	2009
6	7-19, 1-(13)	2008
7	8-19, 1-(14)	2007
8	9-19, 1-(15)	2006
9	10-19, 1-(16)	2005
10	11-19, 1-(17)	2004
11	12-19, 1-(18)	2003
12	13-20, 1-(18)	2003
13	14-19, 1-19, (1)	2001, 2020
14	15-20, 1-19, (1)	2001, 2020
15	16-19, 1-19, 1-(3)	1999, 2018
16	17-19, 1-19, 1-(4)	1998, 2017
17	18-19, 1-19, 1-(5)	1997, 2016
18	19, 1-19, 1-(6)	1996, 2015
19	1-19, 1-(7)	1995, 2014
20	1-19, 1-(7)	1995, 2014

4.3.3 Policy Scenarios

The scenarios considered in addition to the baseline model are described below, and their results are given in Section 4.4.

Fuel Tax The first policy change we choose to investigate is an increase in the cost to the consumer of gasoline. One mechanism that policymakers have for implementing such a change is to increase the tax collected per gallon. As suggested in [21], we implement at the start of year 2001 a gasoline tax increased to the level charged in the United Kingdom in 2001 [16], which represents an increase of \$3.06/gallon (year 2000 dollars).

Hybrid Only Another policy that we analyze with our population of vehicles is the widespread availability of a vehicle with significantly improved fuel economy. Data for the profile of a hybrid vehicle to be introduced to the consumers in the year 2004 are given in Table 4.2. It should be noted that the Toyota Prius, on

which our hybrid vehicle profile is based, is a smaller vehicle than the generic, domestic, 5-passenger sedan used in the baseline scenario, and hence we do not offer consumers an identical challenger. This shortcoming is addressed in the Multiple Challengers scenario.

Fuel Tax + Hybrid We examine a scenario in which our two policies coincide: availability of an efficient hybrid vehicle in 2004, and a significantly increased gasoline tax in 2001.

Multiple Challengers This scenario is similar to the *Hybrid Only* scenario in that a hybrid vehicle becomes available to consumers in the year 2004. In this case, however, a fixed percentage of consumers have the choice to buy a traditional or hybrid vehicle whereas above they were restricted to considering only the hybrid model. As each household is initialized, they are assigned a category which allows a 20% chance of their being allowed to choose between hybrid and traditional vehicles from 2004 on. The remaining 80% of the households only have the traditional, baseline vehicle to consider in their replacement decisions. This limitation on the number of hybrids that may be introduced is intended to capture two problems with the Hybrid Only scenario. The first is that the Toyota Prius on which we model our hybrid vehicle is smaller than the mid-size, 5-passenger sedan we use in the baseline case. A consumer will have to be flexible in his or her vehicle functionality to switch to the smaller vehicle, and we recognize this would not happen in all cases. The second cause for limiting the distribution of the hybrid vehicle choice is that currently demand exceeds supply. That is, only a limited number of hybrid vehicles are available for sale each year (an estimated high of 47,000 Prius' in 2004 [2]), and therefore it does

not make sense to allow our entire population of consumers to switch at once.

Alternative Consumer Behavior Now consider the effect on consumers' replacement policies if, instead of minimizing cost, the consumers minimize their $NMHC + NO_x$ emissions. We saw in Section 2.4 that shorter vehicle lives reduce the cumulative $NMHC + NO_x$ emissions as new technology is brought onto the road, removing older, polluting technology. We repeat the baseline multivehicle scenario described in Section 4.3.2 with the following changes, and the results are given in Section 4.4.5.

1. All consumers minimize $NMHC + NO_x$ only
2. We perform this second analysis twice. First, we randomly make 20% of the population minimize $NMHC + NO_x$ only while the remaining 80% of consumers only minimize the economic cost (as in the Baseline scenario). In the second trial, we let 50% of the population minimize $NMHC + NO_x$ only while the remaining half minimize only economic cost.
3. We randomly generate a weighted objective function for each household. The weight for $NMHC + NO_x$ may be up to 50%, and the remaining weight is applied to the economic cost objective.

4.4 Results

As we examine the output of the consumers given the policy scenarios discussed above, we consider two questions of interest. The first of these is: Does consumer vehicle replacement behavior change from that seen in the Baseline scenario? We display the new consumer replacement policies in tables similar to Table 4.1 in order to compare this information. Our second question of interest is: Does the population

Table 4.2: Data for Hybrid Vehicle Profile

Category	Reference Value	Source
Cost		
Purchase	20810	2004 Toyota Prius MSRP
Fuel Economy	56 mpg (2004)	www.fueleconomy.gov
Maintenance		<i>same as baseline</i>
End of Life		<i>same as baseline</i>
Carbon dioxide, CO_2		
Production	2786 kg	[45] pgs. 4-12, 4-13
Use	158 g/mi	[42] Prius lab test emissions
Maintenance		<i>same as baseline</i>
End of Life		<i>same as baseline</i>
Nitrogen Oxides, NO_x		
Production		<i>same as baseline</i>
Use	0.010 g/mi	[42] Prius lab test emissions
Maintenance		<i>same as baseline</i>
End of Life		<i>same as baseline</i>
Non-methane hydrocarbons, $NMHC$		
Production		<i>same as baseline</i>
Use	0.0024 g/mi	[42] Prius lab test emissions
Maintenance		<i>same as baseline</i>
End of Life		<i>same as baseline</i>

profile (for economic or emissions burdens) change from the Baseline scenario? We display this information by plotting the average annual burden per household under the various scenarios. The following results are for each of the policy scenarios described in Section 4.3.3.

4.4.1 Fuel Tax

When the increase in the cost of gasoline due to a new tax policy is introduced in 2001, there is only one change in behavior for consumers driving the baseline vehicle. Those consumers whose vehicle was 2 years old at the end of 1994 now only keep it through age 19, replacing a year earlier so the new vehicle is 9 years old at the end of 2020. This earlier replacement occurs in 2012. Because there is little consumer behavior change due to this policy, the increase in the average annual cost per consumer over the Baseline scenario shown in Figure 4.2 is primarily due directly to the same number of gallons being sold merely at higher prices. In Figure 4.3 we

again see the result of the minimal behavior change in the fact that the average annual burden of $NMHC + NO_x$ varies little from the Baseline scenario case.




Figure 4.2: Average Annual Cost: Fuel Tax vs. Baseline scenarios

4.4.2 Hybrid Only

Now suppose that a hybrid vehicle resembling the Toyota Prius, as described in Table 4.2 is introduced as the only vehicle available for purchase in the year 2004. The consumer behavior changes dramatically, as shown in Table 4.3. The sections of the table set off between double lines show that many of the consumers ended up on the same vehicle replacement schedule once the hybrid was introduced. This causes the always significant purchase costs to occur in the same years for most households, and the resulting average annual cost per household shown in Figure 4.4 contains economically unrealistic peaks corresponding to these popular purchase years. In Figure 4.5, we see a dramatic reduction of $NMHC + NO_x$ emissions as all consumers




Figure 4.3: Average Annual $NMHC + NO_x$ emissions: Fuel Tax vs. Baseline scenarios

eventually replace with the hybrid vehicle.

4.4.3 Fuel Tax + Hybrid

When we introduce the gasoline cost increase in 2001 to this population which replaces with hybrids from 2004, we see a shift in the hybrid vehicle lifetimes. These new policies are shown in Table 4.4. Similarly, we see only moderate shifts in the corresponding profiles of average annual cost and $NMHC + NO_x$ emissions; these are shown in Figures 4.6 and 4.7, respectively.

4.4.4 Multiple Challengers

Now we examine the results of allowing just 20% of the population to be able to replace with a hybrid beginning in 2004. The resulting population cost profile is shown in Figure 4.8, and the population $NMHC + NO_x$ profile is shown in Figure 4.9.

Table 4.3: Consumer behavior from Hybrid Only scenario

Age at end of 1994	Ages of vehicles owned 1995-2020	Replacement Years
1	2-19, 1-(8)	1994, 2013
2	3-19, 1-(9)	2012
3	4-19, 1-(10)	2011
4	5-20, 1-(10)	2011
5	6-14, 1-13, 1-(4)	2004, 2017
6	7-15, 1-13, 1-(4)	2004, 2017
7	8-16, 1-13, 1-(4)	2004, 2017
8	9-17, 1-13, 1-(4)	2004, 2017
9	10-18, 1-13, 1-(4)	2004, 2017
10	11-19, 1-13, 1-(4)	2004, 2017
11	12-20, 1-13, 1-(4)	2004, 2017
12	13-18, 1-12, 1-(8)	2001, 2013
13	14-19, 1-12, 1-(8)	2001, 2013
14	15-19, 1-13, 1-(8)	2000, 2013
15	16-19, 1-14, 1-(8)	1999, 2013
16	17-18, 1-16, 1-(8)	1997, 2013
17	18-19, 1-16, 1-(8)	1997, 2013
18	19, 1-17, 1-(8)	1996, 2013
19	1-18, 1-(8)	1995, 2013
20	1-18, 1-(8)	1995, 2013

Table 4.4: Consumer behavior from Fuel Tax + Hybrid scenario

Age at end of 1994	Ages of vehicles owned 1995-2020	Replacement Years
1	2-17, 1-(10)	1994, 2011
2	3-18, 1-(10)	2011
3 (<i>no change</i>)	4-19, 1-(10)	2011
4 (<i>no change</i>)	5-20, 1-(10)	2011
5	6-14, 1-16, (1)	2004, 2019
6	7-15, 1-16, (1)	2004, 2019
7	8-16, 1-16, (1)	2004, 2019
8	9-17, 1-16, (1)	2004, 2019
9	10-18, 1-16, (1)	2004, 2019
10	11-19, 1-16, (1)	2004, 2019
11	12-20, 1-16, (1)	2004, 2019
12	13-18, 1-10, 1-(10)	2001, 2011
13	14-19, 1-10, 1-(10)	2001, 2011
14	15-18, 1-13, 1-(9)	1999, 2012
15	16-19, 1-13, 1-(9)	1999, 2012
16	17-19, 1-13, 1-(10)	1998, 2011
17	18-19, 1-14, 1-(10)	1997, 2011
18	19, 1-16, 1-(9)	1996, 2012
19	1-16, 1-(10)	1995, 2011
20	1-16, 1-(10)	1995, 2011




Figure 4.4: Average Annual Cost: Hybrid Only vs. Baseline scenarios

In these cases, the resultant population-wide burdens are not greatly affected, and this seems likely to be the way in which a nation-wide population experiences the introduction of such a technology alternative.

4.4.5 Alternative Consumer Behavior

As described above in Section 4.3.3, we consider three cases of Alternative Consumer Behavior, the results for which are given separately.

1. First we observe the impacts of consumers choosing to minimize only their $NMHC + NO_x$ emissions (instead of only economic cost). Except for consumers owning very new vehicles at the end of 1994, the majority of the population makes an immediate decision to move to a policy allowing new vehicle purchases in 1995, 2001 and 2007. The aligning of all consumers is likely due to a sensitivity in the data making these specific choices better than small vari-




Figure 4.5: Average Annual $NMHC + NO_x$ emissions: Hybrid Only vs. Baseline scenarios

ations on these replacement policies, but we do see in Table 4.5 that the older technology is removed from the road much more quickly than in the Baseline consumer population that minimizes cost. The aligning of consumers unfortunately causes the same spikes in average annual cost and emissions that we saw in the Hybrid Only scenario. Figures 4.10 and 4.11 show the changes in the average population profile from the Baseline scenario in which only economic cost was minimized to the Alternative Consumer Behavior scenario 1 in which only $NMHC + NO_x$ emissions are minimized. Despite the purchase-year peaks, we see that the $NMHC + NO_x$ emissions are reduced much earlier than in the Baseline scenario.

2. We had two experiments that mixed consumers with very different objectives. First we let 20% of the consumers minimize only $NMHC + NO_x$ emissions (while




Figure 4.6: Average Annual Cost: Fuel Tax + Hybrid vs. Baseline scenarios

the remaining consumers minimized cost as in the Baseline scenario), and these results are shown in Figures 4.12 and 4.13. In the second case we increased the number of consumers minimizing only $NMHC + NO_x$ emissions to be 50% of the population. The results for this population are give in Figures 4.14 and 4.15. In both of these trials we see moderate economic cost spikes in the same years as the Alternative Consumer Behavior 1 scenario. This is due to the fact that the individuals minimizing $NMHC + NO_x$ switch immediately to that higher-frequency replacement schedule. It seems as though this result paints a realistic picture of the introduction of low-emission technology. The “Green Consumers” who strive only to reduce their emissions will have to pay monetarily for that objective.

3. The final Alternative Consumer Behavior scenario explored allows consumers




Figure 4.7: Average Annual $NMHC + NO_x$ emissions: Fuel Tax + Hybrid vs. Baseline scenarios

to minimized a weighted sum of economic cost and $NMHC + NO_x$ emissions. These results are given in Figures 4.16 and 4.17, respectively. Here we see only small, gradual changes in the cost and $NMHC + NO_x$ burdens, as the magnitude of the economic cost allows it to continue to be favored above the $NMHC + NO_x$ emissions. We do, however, see the effect of individuals trying to consider multiple objectives, and this seems to reflect somewhat the trend of consumer education [2].

4.5 Discussion

We have now seen the impact on consumer behavior and on average annual burdens of a variety of potential changes in an automotive consumer's environment. It is important to note that each consumer in our model minimizes his or her objective




Figure 4.8: Average Annual Cost: Multiple Challengers vs. Baseline scenarios

over the entire 1995 - 2050 timeline. That is, in the Baseline scenario, for example, a consumer minimizes the total cost to be incurred from the present through 2050, rather than just the cost in the next year. To show relative changes in these cumulative, horizon-length burdens, we display the average values for our population of 200 consumers in Table 4.6. Figures 4.18, 4.19, and 4.20 show the relative changes in each objective as the various policies are implemented.

The other measure of change that we discussed in Section 4.4 is the change in consumer replacement policies. One way in which we observe this change is by analyzing the average age of the vehicles in the population in a given year. When replacements become more frequent, this average age tends to decrease. In Table 4.7, we display the maximum, minimum and mean population-wide average vehicle age for each of our policy scenarios. As we might expect, the scenarios in which we had any con-

Table 4.5: Consumer behavior from Alternative Consumer Behavior 1 scenario

Age at end of 1994	Ages of vehicles owned 1995-2020	Replacement Years
1	2-6, 1-7, 1-(14)	2000, 2007
2	3-7, 1-7, 1-(14)	2000, 2007
3	4-6, 1-9, 1-(14)	1998, 2007
4	1-6, 1-6, 1-(14)	1995, 2001, 2007
5	1-6, 1-6, 1-(14)	1995, 2001, 2007
6	1-6, 1-6, 1-(14)	1995, 2001, 2007
7	1-6, 1-6, 1-(14)	1995, 2001, 2007
8	1-6, 1-6, 1-(14)	1995, 2001, 2007
9	1-6, 1-6, 1-(14)	1995, 2001, 2007
10	1-6, 1-6, 1-(14)	1995, 2001, 2007
11	1-6, 1-6, 1-(14)	1995, 2001, 2007
12	1-6, 1-6, 1-(14)	1995, 2001, 2007
13	1-6, 1-6, 1-(14)	1995, 2001, 2007
14	1-6, 1-6, 1-(14)	1995, 2001, 2007
15	1-6, 1-6, 1-(14)	1995, 2001, 2007
16	1-6, 1-6, 1-(14)	1995, 2001, 2007
17	1-6, 1-6, 1-(14)	1995, 2001, 2007
18	1-6, 1-6, 1-(14)	1995, 2001, 2007
19	1-6, 1-6, 1-(14)	1995, 2001, 2007
20	1-6, 1-6, 1-(14)	1995, 2001, 2007

Table 4.6: Cumulative burdens over the 1995 - 2050 horizon

Scenario	Cost (1975\$)	NMHC + NO _x (kg)	CO ₂ (kg)
Baseline	36188.13	1034.96	306710.09
Fuel Tax	46457.11	1033.50	306669.02
Hybrid Only	33844.99	639.44	174907.00
Fuel Tax + Hybrid	40366.72	613.14	171507.23
Multiple Challengers	35877.09	980.90	286990.14
Alt Consumer Behavior 1	46081.19	721.97	316084.40
Alt Consumer Behavior 2 (20%)	38376.29	963.69	308815.65
Alt Consumer Behavior 2 (50%)	41241.47	873.15	311510.97
Alt Consumer Behavior 3	36186.22	1032.10	306879.71




Figure 4.9: Average Annual $NMHC + NO_x$ emissions: Multiple Challengers vs. Baseline scenarios

sumers minimize $NMHC + NO_x$ emissions resulted in lower average population-wide vehicle ages than the scenarios in which economic cost was minimized.

As we were careful to state at the start of this Chapter, it is important to approach computer modeling of this kind of large, complex system with care. We have implemented a number of possible changes to the environment for a population of vehicle owners, and we have observed the results of these changes given our assumptions

Table 4.7: Population-wide age averages for each scenario

Scenario	Maximum	Minimum	Mean
Baseline	11.4	9.1	10.0
Fuel Tax	11.4	9.1	10.0
Hybrid Only	12.2	3.9	8.4
Fuel Tax + Hybrid	12.1	4.4	8.7
Multiple Challengers	11.0	7.7	9.4
Alt Consumer Behavior 1	14.0	1.0	5.7
Alt Consumer Behavior 2 (20%)	11.0	7.6	9.0
Alt Consumer Behavior 2 (50%)	11.8	5.2	7.8
Alt Consumer Behavior 3	11.4	9.2	10.0




Figure 4.10: Average Annual Cost: Alternative Consumer Behavior 1 vs. Baseline scenarios

regarding the consumers' behavior and motivations. There are a couple of model shortcomings that should be noted as the results are presented.

The first area of concern regards the profile of our hybrid vehicle. As discussed, it is modeled after the Toyota Prius, and the fact that the size and performance of the Prius differ from the mid-size, passenger vehicle previously available changes the functionality that we offer to our consumers. Additionally, we report the same categories of emissions for our hybrid vehicle as for the baseline vehicle, but there has been concern raised that the technology employed in the hybrid vehicles differs sufficiently from traditional internal combustion engine vehicles that perhaps additional pollutants (such as lead) should be studied as well [44].

A second opportunity for improvement arises in the Alternative Consumer Behavior 3 scenario. The weighting scheme employed for each consumer is a naive




Figure 4.11: Average Annual $NMHC + NO_x$ emissions: Alternative Consumer Behavior 1 vs. Baseline scenarios

summation of noncommensurate, weighted values. It is possible that an approach which normalizes the individual objectives first could be applied to “level the playing field” for the objectives. This would allow somewhat greater pressure from the $NMHC + NO_x$ objective, for those consumers desiring such an effect.

There are, of course, an endless number of policy alternatives and other external forces which could be evaluated for our population of consumers. These include expansion of the vehicle models studied, acquisition of multiple vehicles by some households, implementation of a Corporate Average Fuel Economy (CAFE) restriction on the vehicles available each year, and a system that allows some consumers to purchase used vehicles from other consumers wishing to upgrade. The possibilities are discussed further in Chapter V, as we see that the multivehicle, multiobjective model of vehicle ownership allows us the flexibility to ask a great variety of ques-




Figure 4.12: Average Annual Cost: Alternative Consumer Behavior 2 (20%) vs. Baseline scenarios.




Figure 4.13: Average Annual $NMHC + NO_x$ emissions: Alternative Consumer Behavior 2 (20%) vs. Baseline scenarios




Figure 4.14: Average Annual Cost: Alternative Consumer Behavior 2 (50%) vs. Baseline scenarios




Figure 4.15: Average Annual $NMHC + NO_x$ emissions: Alternative Consumer Behavior 2 (50%) vs. Baseline scenarios




Figure 4.16: Average Annual Cost: Alternative Consumer Behavior 3 vs. Baseline scenarios




Figure 4.17: Average Annual $NMHC + NO_x$ emissions: Alternative Consumer Behavior 3 vs. Baseline scenarios




Figure 4.18: Cumulative Economic Cost




Figure 4.19: Cumulative *NMHC* + *NO_x* Emissions




Figure 4.20: Cumulative *CO₂* Emissions

CHAPTER V

Conclusions

In this work we explore a new perspective on the asset replacement problem. The simultaneous consideration of noncommensurate and conflicting environmental and economic objectives significantly increases the complexity of determining the optimal service life of an asset, particularly as we do not restrict ourselves to a specific decision-maker. As part of our collaborative work with research partners in the University of Michigan School of Natural Resources and Environment, we were able to build upon a foundation that included traditional operations research techniques in addition to their tools of Life Cycle Analysis. This partnership enhanced the quality of data and insight into energy and emissions modeling to extend our analysis beyond the hypothetical.

In Chapter II we construct a model to explore the optimal life cycle of a passenger vehicle from the perspective of a variety of objectives. The ease of implementation of a single-objective, single-vehicle dynamic program allows us to quickly observe the tensions among several of the objectives. The extension to the development of multiobjective metaheuristics in Chapter III results in a generous set of Pareto-optimal solutions for the vehicle replacement problem, and from these solutions we glean further information regarding trade-offs available to a decision-maker. With-

out restricting ourselves to a weighted objective function, we gain a set of possible replacement policies for comparison. Additionally, the use of three fundamental algorithms upon which our variants are based allows us to compare the key components of each of these. We discover that the relatively small number of parameters required for the NSGA-II [40] variant contributes to its robust performance, for example.

In Chapter IV we expand the scope of our analysis in recognition of the fact that the decisions of the various stakeholders affecting vehicle service lives are not independent. We assemble a population of consumers owning our Primary Use Vehicles, and we observe the changes in their replacement policies and the resulting vehicle population as characteristics of their environment are altered. We see that whereas some policies (a dramatic increase in fuel cost, for example) have little effect on consumer replacement behavior in the baseline case, other changes (such as the sudden availability of an inexpensive, fuel-efficient, alternative vehicle) change the optimal policy even for those consumers minimizing only economic cost.

The multivehicle model of Chapter IV has many possibilities for future extensions. At present, we consider only one baseline vehicle: the domestic, midsize, 5-passenger sedan owned as a Primary Use Vehicle. The development of additional initial vehicle profiles would allow the consideration of new consumer preferences and behaviors. The more variety available in initial circumstances, the larger the ensemble of models that may be drawn upon for insight, and the greater chance the computer models have for aiding in the discovery of robust policies [5].

Another area of extension for the multi-vehicle model is the consideration of different kinds of policies. We do not presently restrict the annual population-wide emissions, for example, and this could be an interesting area for investigation. Applying this type of “budgetary” constraint would involve new and significant modeling chal-

lenges, but the resulting behavior could enhance insight into the repercussions of the population's individuals being required to work together to reduce their emissions.

A possible alternative to the budget constraint that could achieve similar insights would be to assemble consumers on a spatial location-based network. It is possible that consumers could work together, using penalties to encourage highly-polluting neighbors to alter their behaviors. Additionally, if consumers are given the opportunity to change location (and some initial pollution-tolerance level), would we see segregation of vehicle-owners based on this attitude toward/tolerance for pollution?

As we work to enhance our models of these various levels of stakeholders who are called upon to weigh conflicting objectives in their decision-making processes, it seems that another area of application for these kinds of models would be the healthcare industry. In that case we have policymakers or administrators, caregivers, and patients making decisions that must weigh cost, efficiency and quality of care. The decisions that one patient makes may influence the alternatives available to another patient, and these in turn affect and are shaped by the choices made by doctors and administrators.

Our analysis in this work has focused on the replacement of passenger vehicles, but the multiobjective life cycle analysis tools could be used for a variety of assets whose environmental profiles are in tension with the economic characteristics. Household appliances, lawnmowers and fleets of larger vehicles are just a few examples. Once the (nontrivial) Life Cycle Inventory of data is compiled and feasible alternatives are identified, the methods employed here should be adaptable to these new assets in a straightforward manner.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] *Consumer Expenditure Survey*. U.S. Department of Labor Bureau of Labor Statistics, <http://bls.gov/cex/>, 2000.
- [2] At least for now, hybrid vehicles are still a curiosity. *The Buffalo News*, June 01 2004.
- [3] M.A. Abo-Sinna and M.L. Hussein. An algorithm for generating efficient solutions of multiobjective dynamic programming problems. *European Journal of Operational Research*, 80:156–165, 1995.
- [4] S.C. Bankes. Exploratory modeling for policy analysis. *Operations Research*, 41:435–449, 1993.
- [5] S.C. Bankes. Tools and techniques for developing policies for complex and uncertain systems. *Proceedings of the National Academy of Sciences*, 90:7263–7266, 2002.
- [6] J.C. Bean. Genetic Algorithms and Random Keys for Sequencing and Optimization. *ORSA Journal on Computing*, 6:154–160, 1994.
- [7] T. Bui. Building agent-based corporate information systems: An application to telemedicine. *European Journal of Operational Research*, 122:242–257, 2000.
- [8] J.V. Carnahan and D.L. Thurston. Tradeoff Modeling for Product and Manufacturing Process Design for the Environment. *Journal of Industrial Ecology*, 2:79–92, 1998.
- [9] S. Chand and S. Sethi. Planning horizon procedures for machine replacement models with several possible replacement alternatives. *Naval Research Logistics Quarterly*, 29:483–493, 1982.
- [10] V. Chankong and Y.Y. Haimes. *Multiobjective Decision Making: Theory and Methodology*. Elsevier Science Publishing Co., Inc., New York, New York, 1983.
- [11] S. Choo and P.L. Mokhtarian. What type of vehicle do people drive? The role of attitude and lifestyle in influencing vehicle type choice. *Transportation Research Part A*, 38:201–222, 2004.
- [12] C.A.C. Coello. An updated survey of GA-based multiobjective optimization techniques. *ACM Comput. Surveys*, 32:109–143, 2000.
- [13] D. Haag and M. Kaupenjohann. Parameters, prediction, post-normal science and the precautionary principle - a roadmap for modelling for decision-making. *Ecological Modelling*, 144:45–60, 2001.
- [14] L. Davis. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, New York, 1991.
- [15] S.C. Davis. *Transportation Energy Data Book 20*. Center for Transportation Analysis, Oak Ridge National Laboratory, Oak Ridge, Tennessee, 2000.
- [16] S.C. Davis and S.W. Diegel. *Transportation Energy Data Book 22*. Center for Transportation Analysis, Oak Ridge National Laboratory, Oak Ridge, Tennessee, 2002.

- [17] K. Deb. *Multi-Objective Optimization using Evolutionary Algorithms*. John Wiley & Sons, Ltd, Chichester, 2001.
- [18] E.V. Denardo. *Dynamic Programming: Models and Applications*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1992.
- [19] D.F. Jones, S.K. Mirrazavi and M. Tamiz. Multi-objective meta-heuristics: an overview of the current state-of-the-art. *European Journal of Operational Research*, 137:1–9, 2002.
- [20] D.G. Brown, S.E. Page, R. Riolo and W. Rand. Agent-based and analytic modeling to evaluate the effectiveness of greenbelts. *Environmental Modeling & Software*, In Press, 2004.
- [21] D.V. Spitzley, D.E. Grande, T. Gruhl, G.A. Keoleian and J.C. Bean. Automotive Life Cycle Economics and Replacement Intervals. Technical Report CSS04-01, Center for Sustainable Systems, Ann Arbor, Michigan, 2004.
- [22] E. Zitzler, K. Deb and L. Thiele. Comparison of multiobjective algorithms: empirical results. *Evolutionary Computation*, 8:173–195, 2000.
- [23] E. Elliot and L.D. Kiel. A complex systems approach for developing public policy toward terrorism: an agent-based approach. *Chaos, Solitons and Fractals*, 20:63–68, 2004.
- [24] F. Glover and M. Laguna. *Tabu Search*. Kluwer Academic Publishers, Boston, 1997.
- [25] F. Wenstop and K. Seip. Legitimacy and quality of multi-criteria environmental policy analysis: a meta analysis of five MCE studies in Norway. *Journal of Multi-Criteria Decision Analysis*, 10:53–64, 2001.
- [26] C.M. Fonseca and P.J. Fleming. Genetic algorithms for multiobjective optimization: formulation, discussion and generalization. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 416–423, San Mateo, California, 1993.
- [27] C.M. Fonseca and P.J. Fleming. An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary Computation*, 3:1–16, 1995.
- [28] F. Glover. Tabu Search - Part I. *ORSA Journal on Computing*, 1:190–206, 1989.
- [29] F. Glover. Tabu Search - Part II. *ORSA Journal on Computing*, 2:4–32, 1990.
- [30] F. Glover. Tabu search: a tutorial. *Interfaces*, 20:74–94, 1990.
- [31] D.E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1989.
- [32] M.P. Hansen. Tabu search for multiobjective optimization: MOTS. In *The 13th International Conference on Multiple Criteria Decision Making*, University of Cape Town, January 1997.
- [33] J.H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.
- [34] H.S. Otter, A. van der Veen and H.J. de Vriend. ABLOoM: Location behavior, spatial patterns, and agent-based modelling. *Journal of Artificial Societies and Social Simulation*, 4:1–26, 2001.
- [35] J. Horn, N. Nafpliotis and D.E. Goldberg. A niched pareto genetic algorithm for multiobjective optimization. In *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE Congress on Computational Intelligence*, pages 82–87, New Jersey, June 1994. IEEE Press.
- [36] T. Jackson. Fleet characterization data for MOBILE6. Technical Report M6.FLT.007, U.S. Environmental Protection Agency, 1999.

- [37] J.C. Bean, J.R. Lohmann and R.L. Smith. A dynamic infinite horizon replacement economy decision model. *The Engineering Economist*, 30:99–120, 1985.
- [38] J.C. Bean, J.R. Lohmann and R.L. Smith. Equipment Replacement Under Technological Change. *Naval Research Logistics*, 41:117–128, 1994.
- [39] J.D. Sterman. All models are wrong: reflections on becoming a systems scientist. *System Dynamics Review*, 18:501–531, Winter 2002.
- [40] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6:182–197, 2002.
- [41] K.C. Tan, T.H. Lee and E.F. Khor. Evolutionary algorithms for multi-objective optimization: performance assessments and comparisons. *Artificial Intelligence Review*, 17:253–290, 2002.
- [42] K.J. Kelly and A. Rajagopalan. Benchmarking of OEM Hybrid Electric Vehicles at NREL. Technical Report NREL/TP-540-31086, National Renewable Energy Laboratory, Golden, Colorado, 2001.
- [43] H.C. Kim. *Shaping sustainable vehicle fleet conversion policies based on life cycle optimization and risk analysis*. PhD thesis, University of Michigan, Ann Arbor, MI 48109, 2003.
- [44] L.B. Lave and H.L. MacLean. An environmental-economic evaluation of hybrid electric vehicles: Toyota’s Prius vs. its conventional internal combustion engine Corolla. *Transportation Research Part D*, 7:155–162, 2002.
- [45] M.A. Weiss, J.B. Heywood, E.M. Drake, A. Schafer, and F.F. AuYeung. On the Road in 2020: A life-cycle analysis of new automobile technologies. Technical Report MIT EL 00-003, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2000.
- [46] Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag, New York, 1992.
- [47] P. Nijkamo and J. Spronk. *Multiple Criteria Analysis*. Gower Publishing Co., Ltd., Hampshire, England, 1981.
- [48] B.A. Norman. *The random keys genetic algorithm for complex scheduling problems*. PhD thesis, University of Michigan, Ann Arbor, MI 48109, 1995.
- [49] O.B. Kwon and N. Sadeh. Applying case-based reasoning and multi-agent intelligent system to context-aware comparative shopping. *Decision Support Systems*, 37:199–213, 2004.
- [50] R. Fourer, D.M. Gay and B.W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. International Thompson Publishing, Danvers, Massachusetts, 1993.
- [51] R.J. Lempert, M.E. Schlesinger, S.C. Bankes and N.G. Andronova. The impacts of climate variability on near-term policy choices and the value of information. *Climatic Change*, 45:129–161, 2000.
- [52] C. Romero. *Handbook of Critical Issues in Goal Programming*. Pergamon Press, Oxford, 1991.
- [53] J.D. Schaffer. Multiple objective optimization with vector evaluated genetic algorithms. In J.J. Grefenstette, editor, *Proceedings of the First International Conference on Genetic Algorithms*, pages 93–100, Piscataway, New Jersey, 1985. Lawrence Erlbaum.
- [54] L.M. Schmitt. Theory of genetic algorithms. *Theoretical Computer Science*, 259:1–61, 2001.
- [55] S. Sethi and S. Chand. Planning horizon procedures for machine replacement models. *Management Science*, 25:140–151, 1979.

- [56] S.K. Ko, M.H. Oh and D.G. Fontane. Multiobjective analysis of service-water-transmission systems. *Journal of Water Resources Planning and Management*, pages 78–83, March/April 1997.
- [57] N. Srinivas and K. Deb. Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation*, 2:221–248, 1995.
- [58] C.E. Staley. *A History of Economic Thought: From Aristotle to Arrow*. Blackwell Publishers, Cambridge, Massachusetts, 1989.
- [59] G.L. Thompson. Optimal maintenance policy and sale date of a machine. *Management Science*, 14:543–550, 1968.
- [60] Bureau of Transportation Statistics U.S. Department of Transportation. *NHTS 2001 Highlights Report, BTS03-05*. Washington, DC, 2003.
- [61] D.A. Van Veldhuizen and G.B. Lamont. Multiobjective evolutionary algorithms: analyzing the state-of-the-art. *Evolutionary Computation*, 8:125–147, 2000.
- [62] A. Viana and J.P. de Souza. Using metaheuristics in multiobjective resource constrained project scheduling. *European Journal of Operational Research*, 120:359–374, 2000.
- [63] W. E. Walker. Policy Analysis: A Systematic Approach to Supporting Policymaking in the Public Sector. *Journal of Multi-Criteria Decision Analysis*, 9:11–27, 2000.
- [64] F. Wang and T. Shieh. Extension of interactive dynamic programming to multiobjective optimal control problems. *Industrial and Engineering Chemistry Research*, 36:2279–2286, 1997.
- [65] Y. Nakamori and Y. Sawaragi. Complex systems analysis and environmental modeling. *European Journal of Operational Research*, 122:178–189, 2000.
- [66] E. Zitzler and L. Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation*, 3:257–271, 1999.

ABSTRACT

ASSET REPLACEMENT CONSIDERING ENVIRONMENTAL AND ECONOMIC OBJECTIVES

by

Darby E. Grande

Chair: James C. Bean

Asset replacement considering the impact on a single objective is a well-studied problem, but the extension to multiple objectives poses new challenges. We examine the service life of a passenger vehicle from energy, emissions and economic perspectives. The Life Cycle Optimization performed considers burdens accrued by each objective in the materials production, component manufacturing and assembly, use, maintenance and repair, and end-of-life phases of each vehicle's life. We discover that as technology improves, frequent vehicle replacements, which bring newer technologies onto the road sooner, favor the local emissions objective. In contrast, the economic and global-impact emissions objectives are minimized by longer vehicle lives since their manufacturing burdens are larger in proportion to moderate savings in fuel economy gained by anticipated technological improvements.

We analyze the performance of three heuristic algorithms for the identification of

Pareto-optimal solutions to the multiobjective vehicle replacement problem. These are replacement policies in which the value of one objective may not be improved without degrading the value of another objective. We implement two multiobjective genetic algorithms and one tabu search-based algorithm, and we validate the results using the Constraint Method. In doing this, we discover some commonalities across the algorithms, such as the ability to identify the clustering of solutions into groups corresponding to the number of replacements achieved over the fixed time horizon.

Having determined a representative set of Pareto-optimal solutions for single-vehicle replacement, we employ their trade-off information in constructing a multivehicle model that examines the vehicle replacement decisions of consumer groups. As the model aggregates these individual decisions, we investigate the potential impact of national policy changes such as an increase in fuel tax or widespread introduction of new technology as seen in an advanced, lower-emitting vehicle.