1992

A Bayesian approach to sequential assembly experiments

Klaus Lemke
Iowa State University

Order Number 9311508

A Bayesian approach to sequential assembly experiments

Lemke, Klaus, Ph.D.

Iowa State University, 1992

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

A Bayesian approach to

sequential assembly experiments

by

Klaus Lemke

A Dissertation Submitted to the

Graduate Faculty in Partial Fulfillment of the

Requirements for the Degree of

DOCTOR OF PHILOSOPHY

Departments: Statistics
Industrial and Manufacturing Systems Engineering
Co-majors: Statistics
Industrial Engineering

Approved:

Signature was redacted for privacy.

In Charge of Major Work

Signature was redacted for privacy.

For the Major Departments

Signature was redacted for privacy.

For the Graduate College

Iowa State University
Ames, Iowa
1992

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1. INTRODUCTION

Assemblies such as engines or pumps are integrated systems of component parts or "part types." Performance measurements of assembled units, here also referred to as observations or "test results", vary among units. Since variability in performance is lack of quality, our broad goal is to improve the quality of assemblies by eliminating variability.

We are interested in "important" sources of variability in performance of an assembly that are attributable to part types and their interactions. Our objective is to identify the important sources through assembly tests. A test consists of an "experimental action" that exchanges the component parts of an assembly. To improve the performance of assemblies, those component parts that are ultimately found to affect their performance would be the focus of engineering design changes.

An "experiment" in this study consists of several sequentially performed tests. Each test consists of an "experimental action" involving the component parts of an assembly. The question arises how to select an experimental action for the next test depending on the outcomes of previous tests. The purpose of this study is to develop a Bayesian approach to sequential assembly experiments and to compare it to a technique, currently used by practitioners, which uses a fixed series of tests.

Approaches to Sequential Experimentation with Assemblies

Bhote [1] describes a sequential technique for assembly experiments, attributed to Dorian Shainin, that aims at identifying an important source of variability from a large number of sources with a few sequential assembly tests. Shainin uses two assemblies which lie on opposite ends of a performance scale. Both units are disassembled and parts are successively exchanged for all part types. After each exchange of parts, both units are reassembled and performance is measured on the assemblies. In each exchange, previously exchanged parts are restored to their original units. Shainin effectively applies a special kind of a one–factor–at–a–time experimental design to two assemblies. His technique is based on the notion that this type of part exchange may reverse the roles of the units and thereby unveil an important source.

Bhote [1] lists the following prerequisites for Shainin's technique:

- The technique is applicable, primarily, in assembly operations (but also in process–oriented operations, where there are several similar processes or machines), where good and bad units are found.
- The performance (output) must be measurable and repeatable.
- The units must be capable of disassembly and reassembly without a significant change in the original output.
- There must be at least two assemblies or units – one good and one bad.

According to Bhote's description of Shainin's technique, two assemblies are selected from stockpiles of good and bad units. This implies that assemblies have been sorted prior to experimentation.

We consider a scenario where an assembly consists of three part types and where there is no prior sorting of good and bad units. The effects on assembly performance that we consider are those attributable to parts, the assembly operation, and measurement error. We assume there is no more than one important source of variation attributable to parts.

In this study, assemblies consist of one part of each part type. A unit may be disassembled and reassembled with parts exchanged from another assembly as described above. Alternatively, a unit may be reassembled with parts exchanged from part bins, which are supplies of new parts of each type. These different classes of experimental actions are characteristic for the two approaches that we consider. In particular, our Bayesian approach uses one assembly and part bins as supplies of new parts. Any part type may be exchanged repeatedly and there are no order restrictions on the sequence of exchanges. Shainin's approach, on the other hand, uses one–time part swaps between two assemblies which are performed in a predetermined sequence.

Any experimental action incurs a cost which is subtracted from a budget. We require that the remaining budget is non–negative and thus ensure that we ultimately terminate a sequential assembly experiment. When an experiment terminates we evaluate the decision cost that we incur by making a final decision regarding the unknown identity of the important source of variation. Thus, we consider costs explicitly in our approaches to sequential experimentation with assemblies.

Look–Ahead and Swapping Strategies

Our Bayesian approach includes a "look–ahead" step to simulate sequences of additional tests. We compute the total expected costs for each of these sequences and select the first experimental action of the minimum total expected cost sequence as the candidate action for the next test. If the expected decision cost is lower than the current minimum expected decision cost and the action is feasible within the remaining budget, then we perform the next test. If either of the two criteria is not met, then we stop and attempt to identify the important source by making a decision that minimizes the expected decision cost.

The "swapping" strategy begins by finding two extreme assemblies that lie on opposite ends of a performance scale. We sequentially assemble and measure assemblies until two units, labeled "high" and "low", satisfy a statistical criterion to establish a repeatable difference in performance between these two units. Subsequently, in a predetermined sequence and using one–factor–at–a–time experimental actions, as suggested by Shainin, parts are exchanged between the two extreme units which are reassembled and performance is measured on the assembled units. A reversal of the roles of the extreme units occurs when the performance measurement of the low unit exceeds that of the high unit after a part swap. The notion on which we base our decisions regarding the unknown source of variation is that a reversal reveals an important part type of the unknown source.

Each strategy leads to a heuristic for sequential assembly experiments. A heuristic generally consists of a set of rules for exchanging parts, a stopping rule, and a decision rule. Our heuristics are structured to ensure that total experimental cost does not exceed the initial budget so that we must ultimately terminate an experiment.

CHAPTER 2. LITERATURE SEARCH

Lindley [4] notes that the selection of a sequence of experimental actions is an area of application for Bayesian theory. The sequence of events in time begins with the selection of an experimental action which is followed by the performance of a test. The last step is to reach a decision which incurs a cost depending on the true state.

The mathematical analysis of the sequence of events proceeds in reverse time order. For a linear cost function and a fixed set of experimental actions, including a terminal action, Lindley shows that the Bayesian approach to sequential experimentation leads to a dynamic programming problem. He analyzes the case where, at each stage of the experiment, the experimenter has the choice of either taking another observation or of making a terminal decision.

Dynamic programming problems like Lindley's, and the one that would result from attempting full optimization in the present problem, generally defy analytic solution and only special cases of recurrence relationships have been solved analytically (see the references in [4]). Instead of taking a dynamic programming approach, we describe a (sub–optimal) Bayesian heuristic with a short look–ahead horizon for sequential experimentation with assemblies.

CHAPTER 3. RANDOM EFFECTS AND COST MODELS

Random Effects Model

To model the outcome of an assembly test, we consider a linear random effects model for the difference z_t between the actual performance measurement y_t and a known constant expected value of performance μ. Since we consider the case of a three–part assembly, there are seven effects attributable to part types and their interactions. The part main effects α, β, γ, part interactions $\alpha\beta, \alpha\gamma, \beta\gamma$, $\alpha\beta\gamma$, the assembly operation effect δ, and measurement error ε are all modeled as independent Normal random variables. We write

$$z_t = y_t - \mu =$$

$$\alpha_{a(t)} + \beta_{b(t)} + \gamma_{c(t)} + \alpha\beta_{a(t)b(t)} + \alpha\gamma_{a(t)c(t)} + \beta\gamma_{b(t)c(t)} + \alpha\beta\gamma_{a(t)b(t)c(t)} + \delta_{d(t)} + \varepsilon_t$$

where $\tau \sim N(0, \sigma_\tau^2)$ for $\tau = \alpha, \beta, \gamma, \alpha\beta, \alpha\gamma, \beta\gamma, \alpha\beta\gamma, \delta, \varepsilon$.

For test t, the subscripts $a(t)$, $b(t)$, and $c(t)$ indicate a specific part of part type A, B and C, respectively. The subscript $d(t)$ identifies a unique combination of parts for test t.

A standard assumption is that the random effects are pairwise uncorrelated. Scheffé [5] points out that, for a part interaction effect, exchanging a part of a type involved in the interaction effect yields an independent contribution to z_t. This is true regardless of whether a part of any of the remaining part types involved in the interaction effect is exchanged. Scheffé refers to this as "complete independence" among random effects (see [5], p.240f).

We say a random effect attributable to parts is "active", if its variance is positive. We assume that none or exactly one of the seven part effects is active. The remaining part effects have zero variance. In addition, the assembly operation and measurement error effects have positive variances. The variance of z_t is the sum of the variances of the potentially active effect, the assembly operation, and measurement error. These three variances are assumed to be known.

Our problem is to identify the active effect, if there is one, or otherwise determine that none is active. To denote the true state, we use an integer variable i which takes on a value between zero and seven. Zero denotes the state where none of the effects is active and the integers one through seven represent the seven part effects in the order as they appear in our model for z_t. We denote a value of i that we decide on as being the true state by i'.

Cost Model

The cost $c_k(t)$ of an experimental action k for test t is a function of three cost components:

c_m for measurement or remeasurement

c_r for assembly or reassembly and

c_u for using a new part of type A, B or C.

The cost of a new part is the same for all part types and the total experimental cost of a sequence of tests is the sum of the costs of the actions.

We introduce a decision cost function

$$c(i',i) = \begin{cases} 0 & \text{if } i' = i \\ c' & \text{if } i' \neq i \end{cases}$$

where c' is the penalty cost of wrongly deciding on the value i' when $i \neq i'$ is the true value.

If a correct decision is made ($i' = i$), then the decision cost is zero, otherwise a positive penalty cost c' is incurred. Thus, the same penalty cost applies to cases where an effect is declared active when none is truly active, cases where the wrong effect is identified, or cases where we wrongly decide that there is no active effect.

CHAPTER 4. LOOK–AHEAD HEURISTIC

According to the usual random effects model, a vector $\mathbf{z}_p$ of p observations is p–variate Normal with mean $\mathbf{0}$ and a variance–covariance matrix whose structure depends on which effect, if any, is active and on the sequence of experimental actions. We write

$$\mathbf{Z}_p = \begin{pmatrix} z_1 \\ z_2 \\ \vdots \\ z_p \end{pmatrix} \sim \mathrm{N}_p(\mathbf{0}, \Sigma_{pp}(i)).$$

The variance of the observation z_t for test t is

$$Var(z_t) = \begin{cases} \sigma_\delta^2 + \sigma_\varepsilon^2 & \text{if } i = 0 \\ \sigma^2 + \sigma_\delta^2 + \sigma_\varepsilon^2 & \text{if } i > 0 \end{cases}$$

where σ^2 is the variance of the active effect. The variances of the assembly operation and measurement error effects are denoted by σ_δ^2 and σ_ε^2, respectively.

To indicate which part is used in test t we introduce part indicators, $I_X(t)$, which increase by 1 each time another part of type $X = A, B$, or C is used. In addition, we introduce an assembly indicator, $I(t)$, to identify a unique combination of parts for test t. $I(t)$ is incremented by 1 each time a new combination of parts is tested. When we make an initial assembly ($t = 1$), we set the part indicators and the assembly indicator equal to 1.

To express the covariances between two observations, we define, for each part type X, a binary part indicator $J_X(t_1, t_2)$ as the absolute difference between I_X for tests t_1 and t_2. If the parts of a particular type X differ for two tests, then $J_X = 0$, otherwise $J_X = 1$. We have

$$J_X(t_1, t_2) = \begin{cases} 0 & \text{if } |I_X(t_1) - I_X(t_2)| > 0 \\ 1 & \text{if } |I_X(t_1) - I_X(t_2)| = 0. \end{cases}$$

For all $t_1 \neq t_2$ and for all X there is a binary part indicator $J_X(t_1, t_2)$. Similarly, a binary assembly indicator $J(t_1, t_2)$ is defined as the absolute difference between I for tests t_1 and t_2.

$$Cov_0(z_{t_1}, z_{t_2}) = \begin{cases} \sigma_\delta^2 & \text{if } J(t_1, t_2) = 1 \\ 0 & \text{otherwise} \end{cases}$$

$$Cov_1(z_{t_1}, z_{t_2}) = \begin{cases} \sigma^2 + \sigma_\delta^2 & \text{if } J(t_1, t_2) = 1 \\ \sigma^2 & \text{if } J_A(t_1, t_2) = 1 \text{ and } J(t_1, t_2) = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$Cov_2(z_{t_1}, z_{t_2}) = \begin{cases} \sigma^2 + \sigma_\delta^2 & \text{if } J(t_1, t_2) = 1 \\ \sigma^2 & \text{if } J_B(t_1, t_2) = 1 \text{ and } J(t_1, t_2) = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$Cov_3(z_{t_1}, z_{t_2}) = \begin{cases} \sigma^2 + \sigma_\delta^2 & \text{if } J(t_1, t_2) = 1 \\ \sigma^2 & \text{if } J_C(t_1, t_2) = 1 \text{ and } J(t_1, t_2) = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$Cov_4(z_{t_1}, z_{t_2}) = \begin{cases} \sigma^2 + \sigma_\delta^2 & \text{if } J(t_1,t_2) = 1 \\ \sigma^2 & \text{if } J_A(t_1,t_2) \cdot J_B(t_1,t_2) = 1 \text{ and } J(t_1,t_2) = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$Cov_5(z_{t_1}, z_{t_2}) = \begin{cases} \sigma^2 + \sigma_\delta^2 & \text{if } J(t_1,t_2) = 1 \\ \sigma^2 & \text{if } J_A(t_1,t_2) \cdot J_C(t_1,t_2) = 1 \text{ and } J(t_1,t_2) = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$Cov_6(z_{t_1}, z_{t_2}) = \begin{cases} \sigma^2 + \sigma_\delta^2 & \text{if } J(t_1,t_2) = 1 \\ \sigma^2 & \text{if } J_B(t_1,t_2) \cdot J_C(t_1,t_2) = 1 \text{ and } J(t_1,t_2) = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$Cov_7(z_{t_1}, z_{t_2}) = \begin{cases} \sigma^2 + \sigma_\delta^2 & \text{if } J(t_1,t_2) = 1 \\ \sigma^2 & \text{if } J_A(t_1,t_2) \cdot J_B(t_1,t_2) \cdot J_C(t_1,t_2) = 1 \\ & \text{and } J(t_1,t_2) = 0 \\ 0 & \text{otherwise} \end{cases}$$

where $J(t_1, t_2) = 1$ implies that we are using the same combination of parts for tests t_1 and t_2.

Likelihood Function

The likelihood function is the (joint) density of $\mathbf{z}_p$ viewed as a function of i. We write

$$f(\mathbf{z}_p|i) = (2\pi)^{-p/2}|\Sigma_{pp}(i)|^{-1/2}exp\{-\frac{1}{2}\mathbf{z}_p'\Sigma_{pp}^{-1}(i)\mathbf{z}_p\}$$

where $\Sigma_{pp}(i)$ is a variance–covariance matrix of the observations $\mathbf{z}_p$.

Since experimentation proceeds sequentially, we update a variance–covariance matrix by bordering it with a row and column corresponding to a single additional observation. Hemmerle [3] derives an efficient method to compute the inverse of a matrix by bordering. We use his method to compute the inverse of $\Sigma_{(p+1)(p+1)}(i)$, a variance–covariance matrix of $p+1$ observations, for the likelihood function.

To introduce notation, we consider a vector $\mathbf{z}_{p+q}$ of $p+q$ observations. The variance–covariance matrix $\Sigma_{(p+q)(p+q)}(i)$ of the observations $\mathbf{z}_{p+q}$ can be partitioned as follows

$$\Sigma_{(p+q)(p+q)}(i) = \begin{pmatrix} \Sigma_{pp}(i) & \Sigma_{pq}(i) \\ \Sigma_{qp}(i) & \Sigma_{qq}(i) \end{pmatrix}.$$

We take $q = 1$ and let

$$v = Var(z_{p+1}) - \Sigma_{1p}(i)\Sigma_{pp}^{-1}(i)\Sigma_{p1}(i)$$

and

$$Q = v\Sigma_{pp}^{-1}(i) + \Sigma_{pp}^{-1}(i)\Sigma_{p1}(i)\Sigma_{1p}(i)\Sigma_{pp}^{-1}(i)$$

then the inverse of $\Sigma_{(p+1)(p+1)}$ is (see Hemmerle [3])

$$\Sigma_{(p+1)(p+1)}^{-1}(i) = \frac{1}{v}\begin{pmatrix} Q & -\Sigma_{pp}^{-1}(i)\Sigma_{p1}(i) \\ -\Sigma_{1p}(i)\Sigma_{pp}^{-1}(i) & 1 \end{pmatrix}.$$

Also, the determinant of $\Sigma_{(p+1)(p+1)}$ is (see Graybill [2])

$$|\Sigma_{(p+1)(p+1)}(i)| = v|\Sigma_{pp}(i)|.$$

Prior and Posterior Distributions

Our Bayesian approach to sequential assembly experiments uses a discrete prior probability distribution for i, with probabilities p_i between zero and one. To ensure that the posterior distribution can point to the true value of i, we require that the prior probabilities are strictly positive for all values of i. According to Bayes' formula, the posterior distribution for i given the vector $\mathbf{z}_p$ of p observations is

$$g(i|\mathbf{z}_p) = \frac{p_i f(\mathbf{z}_p|i)}{\sum_i p_i f(\mathbf{z}_p|i)}.$$

Based on the prior, the posterior distribution gives the conditional probabilities that i is the true value given the observations $\mathbf{z}_p$. The prior distribution represents the information about i that is available prior to experimentation. We are effectively assuming that no "expert knowledge" about the unknown identity of the important source of variation is available. Instead, we consider a scenario where it is equally likely that none or one of the seven part effects is active. Thus, we assign a prior probability of $\frac{1}{8}$ to each of the possible values of i.

Minimum Posterior Expected Decision Cost

In this section we show how we use the posterior distribution to decide which effect, if any, is active. Given the observations $\mathbf{z}_p$, we determine the posterior mean decision cost for declaring state i' to hold according to the formula

$$h(i'|\mathbf{z}_p) = \sum_i c(i', i) g(i|\mathbf{z}_p)$$

where $c(i', i)$ is the decision cost function and $g(i|\mathbf{z}_p)$ is the posterior distribution for the effect variable i.

When we terminate an experiment we make a decision that minimizes the posterior mean decision cost over all i'. We let

$$h^*(\mathbf{z}_p) = \min_{i'} \; h(i'|\mathbf{z}_p)$$

where $h^*(\mathbf{z}_p)$ denotes the minimum posterior expected decision cost for $\mathbf{z}_p$.

For the decision cost function described above, the decision rule is equivalent to finding a posterior mode. The minimum posterior expected cost decision may be ambiguous because there may be more than one posterior mode. In this case, we choose the smallest value of i' for which $h(i'|\mathbf{z}_p) = h^*(\mathbf{z}_\mathbf{p})$.

Predictive Distribution of q Additional Observations

We proceed to develop the methodology that enables us to select an experimental action based on the outcomes of previous tests. The first milestone is to find the predictive distribution of q additional observations.

Under the assumed random effects model, the predictive distribution of q additional observations $\mathbf{z}^* = (z_{p+1}, z_{p+2}, \cdots, z_{p+q})$ given i and the observations $\mathbf{z}_p$ in hand is q-variate Normal. We have (see page 12 for notation)

$$\mathbf{Z}^*|i, \mathbf{z}_p \sim \mathbf{N}_q(\Sigma_{qp}(i)\Sigma_{pp}^{-1}(i)\mathbf{z}_p, \Sigma_{qq}(i) - \Sigma_{qp}(i)\Sigma_{pp}^{-1}(i)\Sigma_{pq}(i)).$$

Using standard methods, this well-known result enables us to simulate observations for a specific sequence of q additional tests. Table 4.1 lists the allowable experimental actions (and their costs) for this approach.

Table 4.1: Experimental Actions and Costs for the Look–Ahead Heuristic

k	Description	c_k
1	remeasure	c_m
2	reassemble, measure	$c_r + c_m$
3	get part of type A, reassemble, measure	$c_u + c_r + c_m$
4	get part of type B, reassemble, measure	$c_u + c_r + c_m$
5	get part of type C, reassemble, measure	$c_u + c_r + c_m$
6	get parts of types A and B, reassemble, measure	$2c_u + c_r + c_m$
7	get parts of types A and C, reassemble, measure	$2c_u + c_r + c_m$
8	get parts of types B and C, reassemble, measure	$2c_u + c_r + c_m$
9	get parts of types A, B, and C, assemble, measure	$3c_u + c_r + c_m$

Conditional Expected Decision Cost

The next milestone in selecting an action for the next test is to calculate the conditional (posterior) expected decision cost after q additional tests. The first step is to find the conditional expected decision costs after q additional tests given i and p completed tests. The next step is to multiply these costs by the posterior probabilities for i and, lastly, we sum over i. We have

$$E_{\mathbf{Z}^*|\mathbf{z}_p}[h^*(\mathbf{z}_{p+q})]$$

$$= E_{i|\mathbf{z}_p}[E_{\mathbf{Z}^*|i,\mathbf{z}_p}[h^*(\mathbf{z}_{p+q})]]$$

$$= \sum_i [\int \cdots \int h^*(\mathbf{z}_{p+q}) f(\mathbf{z}^*|i,\mathbf{z}_p) dz_{p+1} \cdots dz_{p+q}]\, g(i|\mathbf{z}_p)$$

where $f(\mathbf{z}^*|i,\mathbf{z}_p)$ is the predictive density for q additional observations and $h^*(\mathbf{z}_{p+q})$ is the minimum posterior expected decision cost for $p+q$ observations. We approximate the expectation $E_{\mathbf{Z}^*|i,\mathbf{z}_p}[h^*(\mathbf{z}_{p+q})]$ through Monte–Carlo simulation.

q–Step Look–Ahead

To select an experimental action for the next test, we consider the following rule for stopping and making a decision that minimizes the posterior expected decision cost after q more tests:

Stop at time p if there is no sequence of $q = 1, 2, 3$ allowable experimental actions (see Table 4.1) for which

$$\sum_{t=p+1}^{p+q} c_k(t) + E_{\mathbf{Z}^*|\mathbf{z}_p}[h^*(\mathbf{z}_{p+q})] < h^*(\mathbf{z}_p)$$

and

$$M_p - \sum_{t=p+1}^{p+q} c_k(t) \geq 0$$

where M_p is the budget remaining after p tests, $c_k(t)$ is the cost of experimental action k for test t, and $h^*(\mathbf{z}_{p+q})$ is the minimum posterior expected decision cost for $p + q$ observations.

The procedure for finding a candidate experimental action is to (1) evaluate the conditional expected decision cost for combinations of 1, 2, and 3 additional tests, (2) add total experimental costs in order to obtain the total conditional expected costs after sequences of q additional tests, (3) find a minimum total cost sequence, and (4) perform the first experimental action in the minimizing sequence, provided two conditions are met.

First, the minimum total conditional expected cost for sequences of additional tests (and making a decision about the unknown identity of the important source after these tests) must be less than the current minimum posterior expected decision cost. Secondly, the remaining budget must be sufficient to carry out the proposed experimental action. If either criteria is not met, then we stop and attempt to identify the important source of variation by making a decision that minimizes the current posterior expected decision cost.

Generally, the conditional expected decision cost decreases with the number of additional tests because the posterior distribution tends to become concentrated on a few values (ideally one value) of i when many observations are available. Thus, we like to have a long look–ahead horizon. Currently, however, the longest computationally feasible horizon is 3 additional tests for the 9 allowable experimental actions listed in Table 4.1. The minimum total conditional expected cost may occur for sequences of fewer than three tests. Thus, we compute total conditional expected costs for $\sum_{q=1}^{3} 9^q = 819$ sequences of up to three additional tests.

CHAPTER 5. SWAPPING HEURISTIC

An alternative strategy for sequential experimentation with assemblies is based on the successive exchange of parts between two selected units. We seek two extreme units that are far apart on a performance scale, because this "essentially" (excluding the interferences from the assembly operation and measurement error) implies that the assemblies each carry a part main or interaction effect that differs "substantially" (both in sign and magnitude) from that of the other unit. Since the important effects are part effects, it is conceivable that the roles of the units reverse when we swap parts between the two extreme assemblies.

We list the allowable experimental actions (and their costs) for the swapping heuristic in Table 5.1. Actions $3', 4'$, and $5'$, used by Shainin, are performed pairwise and in series on the two extreme units. Each test may result in a reversal of the roles of the units. We may naively use the number of observed reversals to decide about

Table 5.1: Experimental Actions and Costs for the Swapping Heuristic

k	Description	c_k
2	reassemble, measure	$c_r + c_m$
$3'$	swap part type A, reassemble, measure	$c_r + c_m$
$4'$	swap part type B, restore part type A, reassemble, measure	$c_r + c_m$
$5'$	swap part type C, restore part type B, reassemble, measure	$c_r + c_m$
9	get parts of types A, B, and C, assemble, measure	$3c_u + c_r + c_m$

the unknown identity of the important source. A naive rule is that if we observe no reversal then we decide that none of the effects is active, one reversal points to a part main effect, two reversals point to the corresponding part two–way interaction effect, and a reversal after each swap points to the part three–way interaction as the active effect.

Reversals and Part Interactions

Figure 5.1 depicts "typical" (based on 144 simulated experiments) success frequencies for the naive rule. The variable "class" denotes the number of part types involved in the active effect (see Chapter 6).

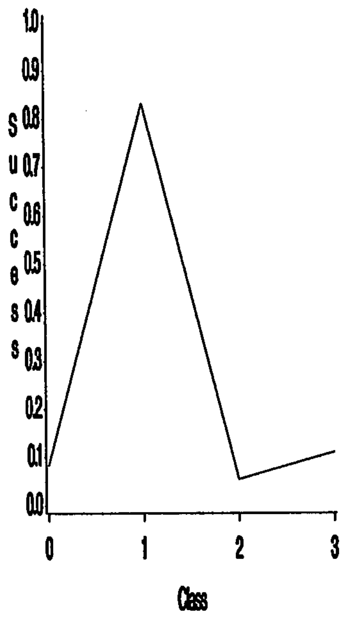

Figure 5.1: Success Frequencies for the Naive Rule

Figure 5.1 shows that only part main effects are successfully identified using this rule. For the following reasons, we do not reliably identify part interactions and cases where none of the effects is active.

Suppose the $\alpha\beta$ part two–way interaction effect is active and let $\alpha\beta_{11}$ and $\alpha\beta_{22}$, where $\alpha\beta_{11} \gg \alpha\beta_{22}$, be the actual effects for the high and low unit, respectively. We assume, for the moment, that there are no interferences from the assembly operation and measurement error and proceed with the first test. Since we are exchanging parts of type A we "draw" two independent effects, $\alpha\beta_{21}$ and $\alpha\beta_{12}$, according to the random effects model.

With probability $\frac{1}{2}$, we have $\alpha\beta_{21} > \alpha\beta_{12}$ so that swapping parts of type A produces no reversal. It is then true that the subsequent swap of parts of type B must yield a reversal. Hence, in this case, we wrongly decide that the β main effect is active. On the other hand, with probability $\frac{1}{2}$, $\alpha\beta_{21} < \alpha\beta_{12}$. Here, swapping parts of type A yields a reversal, while swapping parts of type B does not. In this case, we wrongly decide that the α main effect is active. We conclude that we essentially never make the right decision when a part two–way interaction is active.

A similar rough argument shows that when the part three–way interaction is active, we wrongly decide $\frac{1}{8}$ of the time that none of the effects is active, $\frac{1}{8}$ of the time that the three–way interaction is active, and $\frac{3}{4}$ of the time we wrongly decide on a part main or two–way interaction effect.

When none of the effects is active and with the assembly operation and measurement effects present, we have the same outcomes as for the part three–way interaction because each swap may or may not, with equal probability, result in a reversal of the roles of the two extreme units.

Figure 5.2 depicts the decision tree for a single pass through the sequence of part swaps for part types A, B, and C when the variances for the assembly operation and measurement error effects are small. A swap may result in a reversal (r) or non–reversal (nr). Underneath each r– and nr–node we list the states that are possibly true, where $\emptyset$ denotes the case that none of the effects is active. We observe 0, 1, or 2 reversals. We will actually never observe three reversals because we do not swap parts of type C when we observe a reversal for each of the first two tests. For any number of observed reversals, the list of possible states includes $\emptyset$ and $\alpha\beta\gamma$. When we observe one reversal, the part main and all part two–way interaction effects for which the reversal occurs are added to the above list of possible states. Since we are always left with several possible states after one pass through a series of part swaps, we conclude that we need more than one pass to arrive at an unambiguous decision regarding the unknown identity of the important source. Since we exhausted all part types in the first pass, each additional pass will need to begin by finding two new extreme units.

We propose the following decision rule which is based on two passes. When we observe no reversal in a pass, we decide that none of the effects is active. When two reversals occur (we omit the swap for part type C when we observe reversals for part types A and B) we decide that the three–way interaction effect is active.

We start a second pass when the first pass yields exactly one reversal for some part type. If the second pass also results in a single reversal, we compare the part types for which the reversals occurred. If the part type is the same, then we decide that the corresponding part main effect is active. Otherwise, when the part types differ, we decide that the corresponding part two–way interaction effect is active.

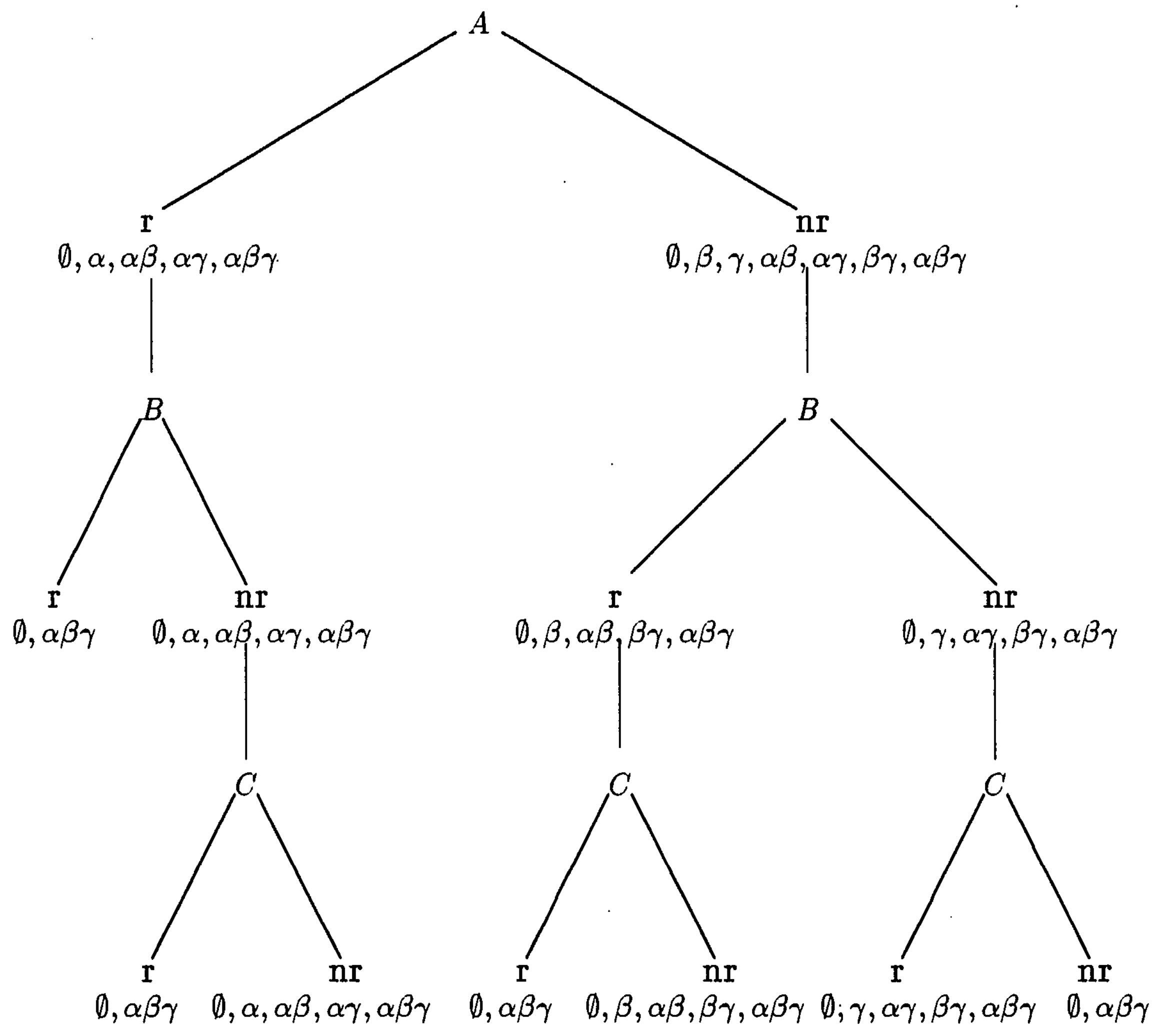

Figure 5.2: Decision Tree for the Swapping Heuristic

Repeatability Ratio

At the beginning of a pass, we assemble two units and rank their performance measurements. These observations are labeled Low_1 and $High_1$, respectively. We reassemble and measure each unit to obtain second measurements, Low_2 and $High_2$, respectively.

Dividing the difference of the mean test results for the low and high units by the square root of the sum of the variances over 2 gives a "repeatability ratio" R. Using a well-known result for the comparison of two means, R has a t–distribution with 2 degrees of freedom under the null hypothesis that the difference in the mean responses for the two units is non–positive.

If we let

$$D = (High_1 + High_2)/2 - (Low_1 + Low_2)/2$$

and

$$d = (High_1 - High_2)^2/2 + (Low_1 - Low_2)^2/2$$

then

$$R = \frac{D}{\sqrt{d/2}} \sim t_2$$

under H_0: $E(D) \leq 0$.

If we do not reject H_0, then we continue and assemble another unit. If its performance is either higher or lower than those of the current extreme units, then we retain this unit and form a new repeatability ratio that includes the original and a remeasurement test result for this unit. We repeat the one–sided t–test. If we reject H_0, then we conclude that we indeed have two extreme units in hand and proceed with the series of part swaps.

We potentially make new assemblies until we reach a lower bound on the remaining budget. For the first pass, this bound is twice the dollar amount required to carry out the series of three part swaps on two assemblies plus the amount required to make two new assemblies for a subsequent pass. If we decide to perform a second pass, then the bound becomes the dollar amount required to carry out the series of three part swaps.

CHAPTER 6. METHODS

Parameters

An "initial" budget represents how much money we are willing to give up in order to be able to identify an important source of variability through subsequent assembly tests. Intuitively, the initial budget does not exceed c', the penalty cost of wrong identification, which is the maximum loss that could be incurred by making a wrong decision. We wish to weigh the experimental effort and the cost of the final decision equally and thus set the penalty cost of wrong identification equal to the initial budget. By setting both these quantities equal to 1,000 dollars we can have a cost structure where the costs of experimental actions are small relative to the initial budget. At the same time, since our heuristics ensure that total experimental cost does not exceed the initial budget, we know that we terminate an experiment after some moderate number of tests.

The look–ahead heuristic uses five Monte–Carlo trials to approximate the conditional expected decision costs for combinations of 1, 2, and 3 experimental actions. The number of trials is kept small because each additional Monte–Carlo trial results in approximately 124,000 additional matrix multiplications for a complete look–ahead step. Finally, our critical value for the repeatability ratio R of the swapping heuristic is $t_{.95,2} = 2.920$.

Factors

The heuristics are "symmetric" in the sense that, as we stay within the part main or part two–way interaction effects, it is unimportant for the success or failure of our heuristics which part main or two–way interaction effect is active. Thus, to avoid redundancies in our calculations and discussion we group states of the same order together into classes. We let the factor "class" denote the number of part types involved in the true active effect. Level 0 of class contains the case where none of the effects is active. The three part main effects α, β, and γ are grouped together into level 1. Similarly, the three part two–way interactions form level 2. Finally, level 3 of class contains the three–way interaction effect. In our simulations, each of the main and two–way interaction effects is selected once to be the active effect. We replicate the remaining levels of class three times.

We use the factor "expense" to denote the costs of experimental actions and consider two levels that we choose to call levels 1 and 3. At level 1 of expense, the cost of a measurement, c_m, is 1 dollar, the reassembly cost, c_r, is 4 dollars, and the cost of using a new part, c_u, is 10 dollars. Level 3 of expense represents a three–fold increase of every cost component. Our cost structure emphasizes the cost of using a new part over the cost of reassembly and measurement.

The variance of an observation z_t when an effect is active, $\sigma^2 + \sigma_\delta^2 + \sigma_\varepsilon^2$, becomes the factor "total variance" in our simulation study (see the formula for $Var(z_t)$ on page 9 in Chapter 4). The levels of total variance are 20 and 80 with $\sigma_\varepsilon^2 = 1$ and 4, respectively. In addition, we define the factor "correlation" to be $\frac{\sigma^2}{\sigma^2+\sigma_\delta^2+\sigma_\varepsilon^2}$ which is an important correlation between two measurements (see the formulas for

$Cov_i(z_{t_1}, z_{t_2})$ on pages 10 and 11 in Chapter 4). This factor can be interpreted as the relative contribution of the active effect to total variance. In our study, the levels of correlation are .5, .7, and .9. Together, the factors total variance and correlation define a variance structure $(\sigma^2, \sigma_\delta^2, \sigma_\varepsilon^2)$. The vectors of variances in our study are (10,9,1), (14,5,1), (18,1,1), (40,36,4), (56,20,4), and (72,4,4).

In summary, we have 3 replications for class with 4 levels, correlation with 3 levels, expense with 2 levels, and total variance with 2 levels. Thus, we simulate $3 \times 4 \times 3 \times 2 \times 2 = 144$ experiments for each heuristic.

Criteria

We consider the following criteria to evaluate the proposed heuristics:

1. frequency of success
2. average number of experimental actions
3. average number of parts used before termination of experimentation and
4. average total experimental cost plus expected decision cost.

The frequency of success approximates the success rate of a heuristic. Total cost becomes large when the total experimental cost or the expected cost for the final decision is large. However, total cost cannot exceed 2,000 dollars for any single experiment, since the initial budget and the penalty cost of wrong identification are 1,000 dollars each.

Analysis

We analyze the results of our simulations in the following ways:

Cross–Tabulation We use a two–way frequency table for the identified effect and the active effect to detect associations between final decisions and true states.

Analysis of Variance We assume a full factorial model using the factors class, correlation, expense, and total variance for a fixed factor univariate analysis of variance of our criteria. The model and error degrees of freedom are 47 and 96, respectively. To us, a p–value less than .01 indicates that a factor or factor combination is important.

Graphical Analysis For the important factors and factor combinations from the analysis of variance, we graphically analyze the observed mean responses for the criteria (see the appendices for graphs).

Main Effect Estimates We compute differences between individual factor level means and a grand mean for each of the criteria. To support our graphical analysis, we construct confidence intervals for linear combinations of main effect estimates.

CHAPTER 7. RESULTS

Look–Ahead Heuristic

Table 7.1 contains the two–way frequencies for the identified and the active effect. This table shows that between 50 and 75 percent of the experiments are successful when one of the effects is active. Overall, two out of three experiments are successful for the look–ahead heuristic.

Table 7.1: Two–way Frequencies for the Look–Ahead Heuristic

Identified	Active Effect								
Effect	None	α	β	γ	$\alpha\beta$	$\alpha\gamma$	$\beta\gamma$	$\alpha\beta\gamma$	Total
None	31	4	1	3	2	2	3	4	50
α	1	6	2	1	0	1	0	3	14
β	1	1	8	1	1	0	0	1	13
γ	0	0	0	7	0	0	0	1	8
$\alpha\beta$	2	1	0	0	9	0	0	1	13
$\alpha\gamma$	0	0	0	0	0	9	1	3	13
$\beta\gamma$	0	0	1	0	0	0	8	0	9
$\alpha\beta\gamma$	1	0	0	0	0	0	0	23	24
Total	36	12	12	12	12	12	12	36	144

In 44 percent of all unsuccessful experiments in Table 7.1 we wrongly decide that none of the effects is active. This is higher than what we expect to find (14 percent) if all states other than the true state were equally likely (wrong) choices.

Using the criterion that $p < .01$ indicates significance, we see from Table 7.2

that the factors class, correlation, and expense affect the success frequency and other criteria. We also notice the significance of the interaction between correlation and expense for the number of actions taken and the number of parts used.

Figure 7.1 contains two separate graphs of the success frequency for different levels of expense. Each graph contains two separate curves corresponding to the two levels of total variance. This figure shows that the success frequency is higher at the low level of total variance than at the high level of total variance.

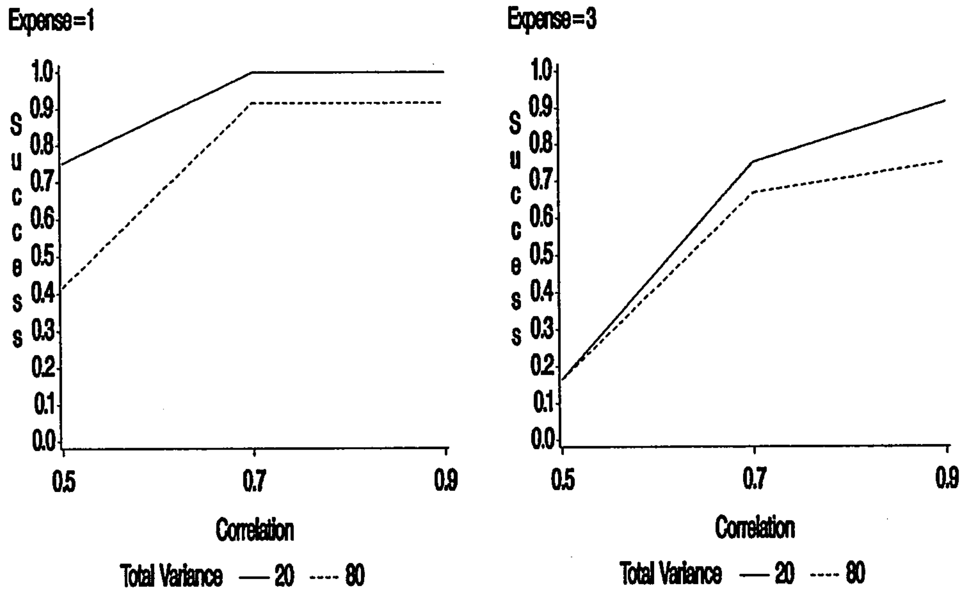

Figure 7.1: Success Frequencies for Two Levels of Expense

More importantly, the success frequency drops below .2 when correlation is .5 and expense is at the high level. The reason for this lack of performance is that 46 percent of the experiments with this factor level combination terminate immediately after the initial assembly has been made and only one in six experiments uses more

Table 7.2: p–Values from Analyses of Variance on Four Criteria for the Look–Ahead Heuristic

Factor Combination	Actions	Parts	Costs	Success
Class (CLA)	.0834	.0029	.0405	.0077
Correlation (COR)	.0001	.0001	.0001	.0001
Expense (EXP)	.0001	.0001	.0001	.0001
Total Variance (VAR)	.2255	.1237	.1474	.0365
CLA*COR	.0758	.0115	.0487	.3111
CLA*EXP	.8674	.3547	.8609	.1463
CLA*VAR	.4927	.1236	.1236	.2509
COR*EXP	.0001	.0001	.2063	.1336
COR*VAR	.3621	.4721	.8976	.8467
EXP*VAR	.3626	.2408	.8867	.4812
CLA*COR*EXP	.8124	.7353	.0231	.4672
CLA*COR*VAR	.9060	.6495	.9823	.2748
CLA*EXP*VAR	.2245	.1464	.7805	.5860
COR*EXP*VAR	.5775	.4916	.5817	.3158
CLA*COR*EXP*VAR	.4380	.3755	.9495	.3725
R^2	.650	.765	.837	.602
$\sqrt{MSE}$	7.652	4.024	143.124	.354

than five parts (three parts for the initial assembly plus one or two parts for additional tests). We find that in 46 percent of the cases of immediate termination we decide that none of the effects is active. The reason for this is the following.

Suppose an initial assembly has been measured and we terminate the experiment. We have an observation z_1 in hand whose variance depends on whether none or one of the effects is active (see the formula for $Var(z_t)$ on page 9 in Chapter 4). Let the posterior probability be g_0 that none of the effects is active. The probability that one of the factorial effects is active, denoted by g_1, is the same for all seven factorial effects. If z_1 is small in absolute value, then $g_0 > g_1$ and we decide that none of the effects is active.

Figure A.1 in Appendix A shows that we are more likely to encounter a successful experiment when none of the effects is active than when there is an active effect. To support this finding, we use the main effect estimates in Table 7.3 to construct a confidence interval for the comparison of the success frequency when class is zero with the success frequencies when class is positive. We have

$$3(.160) - (-.118 + .021 - .062) \pm t_{.975,96}\sqrt{MSE}\sqrt{\frac{12}{36}} = .639 \pm .406$$

where $\sqrt{MSE} = .354$ (see Table 7.2) and $t_{.975,96} = 1.985$. We conclude that the success rate is significantly higher when class is zero than when class is positive.

Since an experimental action uses either 0, 1, 2, or 3 new parts, we have four classes of experimental actions. For example, actions 1 and 2 in Table 4.1 use no new parts, whereas actions 3, 4, and 5 each use one new part. We express the proportion of each class of actions to the total number of actions as a percentage.

Figure 7.2 depicts average percentages of classes of experimental actions for combinations of correlation and expense levels. This figure shows that the percentage

Table 7.3: Main Effect Estimates for the Look–Ahead Heuristic

Factor	Level	Actions	Parts	Costs	Success
Class	0	– 2.57	– 1.52	31.66	.160
	1	1.26	1.97	33.30	– .118
	2	– .38	– .69	– 50.85	.021
	3	1.68	.23	– 14.09	– .062
Correlation	.5	.37	– 1.45	187.58	– .326
	.7	3.41	2.05	28.27	.132
	.9	– 3.80	– .60	– 215.84	.195
Expense	1	5.79	4.50	– 187.33	.132
	3	– 5.79	– 4.50	187.33	– .132
Total Variance	20	– .78	– .52	– 17.42	.062
	80	.78	.52	17.42	– .062
Grand Mean		14.63	11.97	585.33	.701
Standard Error		.88	.57	24.24	.038

of the type of action that uses no new parts increases as correlation decreases. The "mix" of experimental actions changes with the level of correlation because more remeasurement and reassembly experiments are carried out.

Examining the sequence of tests for experiments run during the software development stage (not under the conditions laid out in Chapter 6) we observed that experimental actions that do not involve important part types are chosen for the next test. For example, if the α main effect is active, then the heuristic tends to exchange parts of types B, and C. The reason for this may be that strong correlations among observations provide valuable information for the likelihood function.

We also observe that if, for example, the α main and the $\alpha\beta$ two–way interaction effect both have a large posterior probability of being active, then the action which

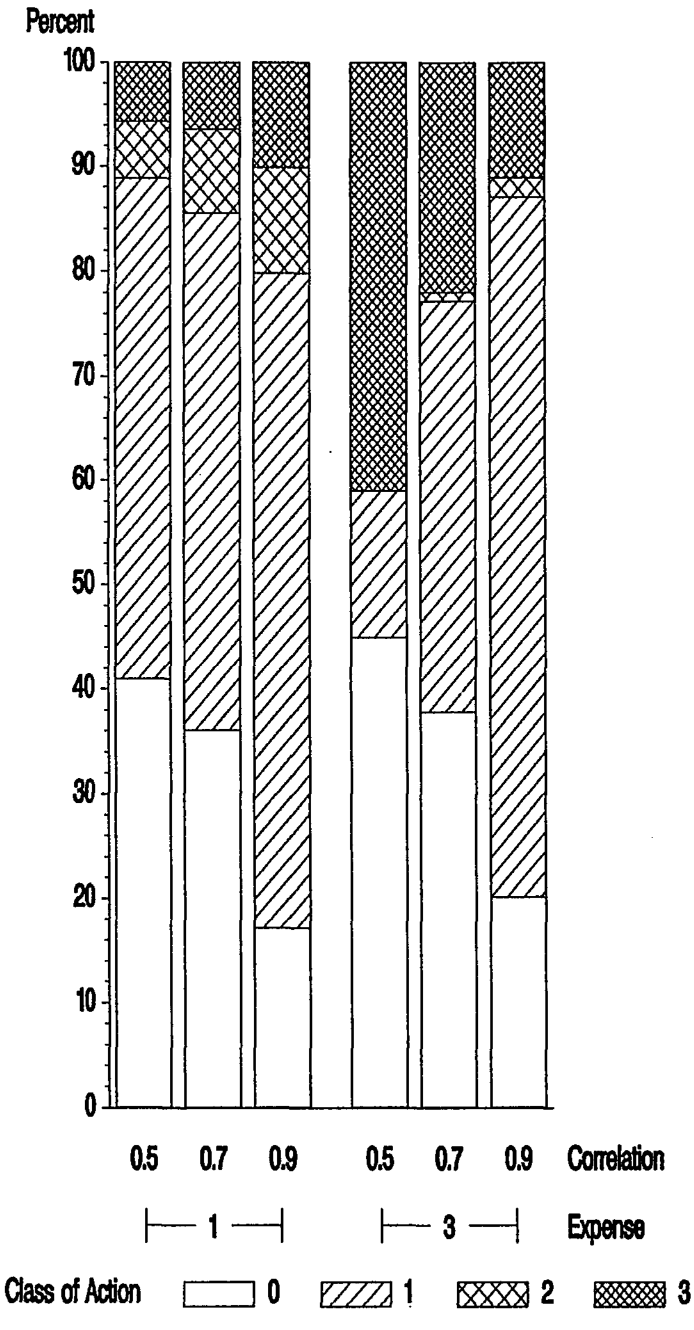

Figure 7.2: Percentages for Classes of Experimental Actions

exchanges parts of type B ($k = 4$ in Table 4.1), is a frequent choice for the next test. If the two most recent test results correlate highly, then part type B is not likely to be involved in the active effect and only the posterior probability for the α main effect remains large after the test.

Swapping Heuristic

Table 7.4 contains the two–way frequencies for the identified and the active effect.

Table 7.4: Two–way Frequencies for the Swapping Heuristic

Identified	Active Effect								
Effect	None	α	β	γ	$\alpha\beta$	$\alpha\gamma$	$\beta\gamma$	$\alpha\beta\gamma$	Total
None	3	1	2	1	4	3	1	6	21
α	1	4	0	0	2	1	0	2	10
β	0	0	7	1	0	0	2	1	11
γ	2	0	0	6	0	1	1	1	11
$\alpha\beta$	3	1	0	0	0	0	0	0	4
$\alpha\gamma$	1	1	0	0	0	2	0	1	5
$\beta\gamma$	0	1	0	1	0	0	3	1	6
$\alpha\beta\gamma$	26	4	3	3	6	5	5	24	76
Total	36	12	12	12	12	12	12	36	144

This table shows that the swapping heuristic, using our proposed decision rule, has little success in identifying a part two–way interaction (14 percent) and only moderate success in identifying a part main effect (47 percent) or the part three–way interaction (67 percent). Overall, one in three experiments is successful for the swapping heuristic.

In 46 percent of all unsuccessful experiments we wrongly decide that the three–way interaction is active. This is higher than what we expect to find (14 percent) if all states other than the true state were equally likely (wrong) choices.

Table 7.5: p–Values from Analyses of Variance on Four Criteria for the Swapping Heuristic

Factor Combination	Actions	Parts	Costs	Success
Class (CLA)	.0001	.0001	.0001	.0001
Correlation (COR)	.0093	.0018	.0001	.0010
Expense (EXP)	.0001	.0001	.0001	.1299
Total Variance (VAR)	.6181	.5842	.5955	.2780
CLA*COR	.0027	.0008	.0073	.1458
CLA*EXP	.0001	.0001	.0001	.0021
CLA*VAR	.4689	.2817	.4717	.4916
COR*EXP	.1394	.1102	.3847	.5406
COR*VAR	.4953	.3741	.8790	.9535
EXP*VAR	.9404	.9418	.4426	.5143
CLA*COR*EXP	.0691	.0172	.2275	.2861
CLA*COR*VAR	.4023	.3929	.4681	.5650
CLA*EXP*VAR	.1981	.1554	.4905	.4916
COR*EXP*VAR	.3957	.4290	.5998	.6527
CLA*COR*EXP*VAR	.2864	.1585	.1991	.5650
R^2	.644	.724	.685	.567
$\sqrt{MSE}$	6.665	13.657	433.399	.382

Using the criterion that $p < .01$ indicates significance, we see from Table 7.5 that the factors class and correlation as well as the interaction between class and expense are important for the success frequency.

The interaction between class and expense is significant in Table 7.5 because we have a low success frequency for part main effects when expense is high and a high success frequency for part main effects when expense is low, while the success frequencies for all other levels of class do not change with the level of expense.

Table 7.6 presents main effect estimates that may be used to construct confidence intervals for comparisons among factor level means.

Table 7.6: Main Effect Estimates for the Swapping Heuristic

Factor	Level	Actions	Parts	Costs	Success
Class	0	4.32	14.00	409.18	- .257
	1	.04	- 3.50	- 136.42	.132
	2	- 1.29	- 4.17	145.28	- .201
	3	- 3.08	- 6.33	- 418.02	.327
Correlation	.5	2.37	5.75	270.83	- .173
	.7	- .63	- 1.63	- 93.02	.077
	.9	- 1.75	- 4.13	- 177.82	.098
Expense	1	3.72	9.21	- 184.16	.048
	3	- 3.72	- 9.21	184.16	- .048
Total Variance	20	.28	.63	- 19.23	.034
	80	- .28	- .63	19.23	- .034
Grand Mean		17.46	24.75	1262.92	.340
Standard Error		.76	1.78	52.61	.040

Comparison of Look–Ahead and Swapping Heuristic

Figure 7.3 depicts the mean responses for the criteria and the important main factors class, correlation, and expense. Each graph contains separate curves for the look–ahead (LA) and swapping (S) heuristics. This figure shows that the look–ahead heuristic has higher success frequencies, takes fewer actions, uses fewer parts, and incurs a lower total cost than the swapping heuristic. (Fisher's least significant differences for pairwise comparisons of grand means for the criteria "Actions", "Parts", "Costs", and "Success" are 1.67, 2.34, 75.02, and .086, respectively.)

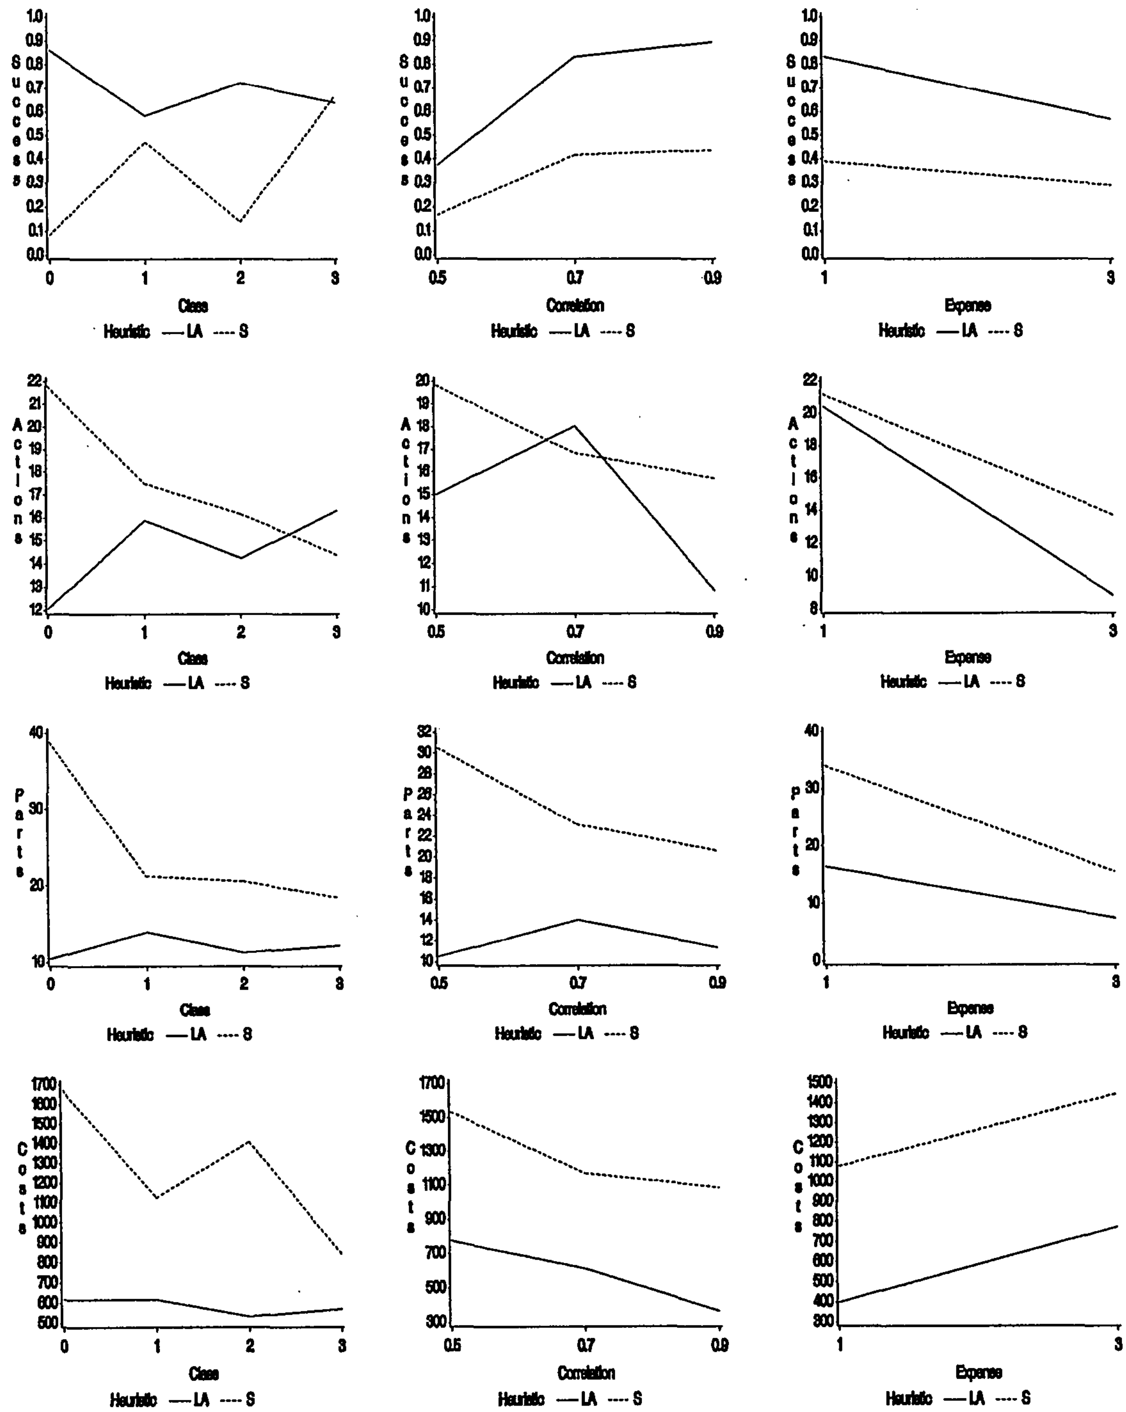

Figure 7.3: Comparison of Look-Ahead and Swapping Heuristic

CHAPTER 8. DISCUSSION AND CONCLUSIONS

Assuming a random effects model for the outcomes of assembly tests, we argued that a naive decision rule for the swapping strategy will not lead to a correct decision when a part 2–way interaction effect is active. The reason is that, in a single pass through a series of part swaps, we essentially observe a reversal for either of the part types involved in a part 2–way interaction effect, but not both. Consequently, we are not allowed to eliminate interaction effects involving a part type for which a reversal has not occurred after a part swap. Similarly, the success rate is essentially $\frac{1}{8}$ when none of the effects is active or when the part 3–way interaction effect is active.

We considered a swapping heuristic that allows two passes through a series of one–factor–at–a–time part swaps. A single pass may show 0, 1, or 2 reversals. When there are two reversals, we decide that the part 3–way interaction effect is active because it is the only remaining possible state. In all other cases, several possibilities remain. To simplify the decision making (and to reduce the chance of deciding on the part 3–way interaction), we exclude the possibility that the part 3–way interaction is active when fewer than two reversals occur. We are left with the decision that none of the effects is active when no reversal occurs. When one reversal occurs, the part main effect or any of the two part 2–way interactions involving the part type for which the single reversal occurred may be the true source. In this case, we find two new extreme assemblies and pass a second time through the series of part swaps. Again, there may be either 0, 1, or 2 reversals. The cases of no reversal and 2 reversals are treated as above. If we observe a single reversal for the same part type in both passes, we pick the part main effect from the list of possibilities. Otherwise, we pick a part 2-way interaction effect corresponding to the two part types for which we noted a reversal in two separate passes through the series of part swaps.

Our Bayesian approach uses a set of experimental actions that is larger than that of the swapping heuristic. Experimental actions may be combined in any conceivable way to form a sequence of tests. This approach includes a look-ahead step in which we evaluate the merits of all possible sequences of up to three additional tests to select an experimental action for the next test. The measure of merit is the sum of the total experimental cost and the conditional expected cost of making a decision regarding the unknown identity of the important source at the end of a particular sequence of additional tests. The sequence with the minimum total conditional expected cost has the greatest merit and we select the first experimental action in this sequence as a candidate action for the next test.

We use Bayes' rule and the set of available test results (which may include additional test results from the look-ahead) to compute the posterior (conditional) probabilities that a particular effect, if any, is active. An expected cost is associated with each decision regarding the unknown identity of the important source and we compare the current expected decision cost to the minimum total conditional expected cost sequence of additional tests to decide if we carry out the next test. We stop and use the results of the completed tests to make a final decision when the current expected decision cost is the smaller of the above two costs.

Our simulations show that the success rate, number of experimental actions, usage of parts, and total cost depend on whether none or one of the effects is active, the variance structure, and the cost structure. Overall, the look–ahead heuristic has a significantly higher success rate and takes fewer experimental actions, uses fewer parts, and incurs lower total cost than the swapping heuristic.

The swapping heuristic did not perform well in our simulations. In particular, its success rate is low when none of the effects or when a part 2–way interaction effect is active. Only one in three experiments ends successfully for the swapping heuristic.

The look–ahead heuristic, on the other hand, is capable of correctly identifying any of the effects or that none of the effects is active. It does so with great precision when there is little interference from the assembly operation and measurement error and when the costs of experimental actions are low. However, the look–ahead heuristic "stalls" when the noise from the assembly operation is substantial and, simultaneously, experimental costs are large relative to the budget (see Figure 7.1). "Stalling" refers to the termination of an experiment immediately after we assemble the initial unit. We observe that stalling is frequently accompanied by a wrong decision that none of the effects is active. Thus, stalling is undesirable. The reason for early termination is that the conditional expected decision cost remains high after 3 additional tests when there is strong interference from the assembly operation and measurement error. After the total experimental costs have been added, the total conditional expected decision cost is not less than the expected decision cost in this case. Thus, we terminate experiments early because the length of the look–ahead horizon is too short to include enough additional tests that could lower the expected decision cost.

When experimental costs are low relative to the initial budget, then the look-ahead generally "adapts" to an increased level of interference from the assembly operation and measurement error by performing a large number of reassembly tests that use no new parts (see Figure 7.2). Adapting is desirable because it appears to be leading to high success frequencies when correlation is .7 or .9 (see Figure 7.1).

In conclusion, the look-ahead strategy dominates the swapping strategy in that it is more successful, uses fewer parts, and is less costly (in terms of total cost) than the swapping heuristic. The reason is that the look-ahead heuristic is flexible, adaptive, and capable of reliably identifying all states. The look-ahead horizon would have to be longer than three experiments to avoid "stalling" of the heuristic when there is strong interference from the assembly operation and measurement error and, at the same time, experimental actions are expensive relative to the initial budget.

In this study, we concentrate on identifying the true active effect, if there is one, in the case of a three-part assembly where variances are assumed to be known. We rarely will be able to make this assumption in practice. Another limitation of our Bayesian approach is that computations for the q-step look-ahead become prohibitive when we consider studying the impact of more than three part types on performance of assemblies.

BIBLIOGRAPHY

[1] Bhote, K.R., 1988. *World Class Quality: Design of experiments made easier, more cost effective than SPC*, American Management Association, New York

[2] Graybill, F.A., 1983. *Matrices with Applications in Statistics.* Wadsworth, Belmont, California

[3] Hemmerle, W.J., 1967. *Statistical Computations on a Digital Computer.* Blaisdell, Waltham, Massachusetts

[4] Lindley, D.V., 1972. *Bayesian Statistics, A Review.* Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania

[5] Scheffé, H., 1963. *The Analysis of Variance.* Wiley, New York

APPENDIX A. RESULTS FOR THE LOOK–AHEAD HEURISTIC

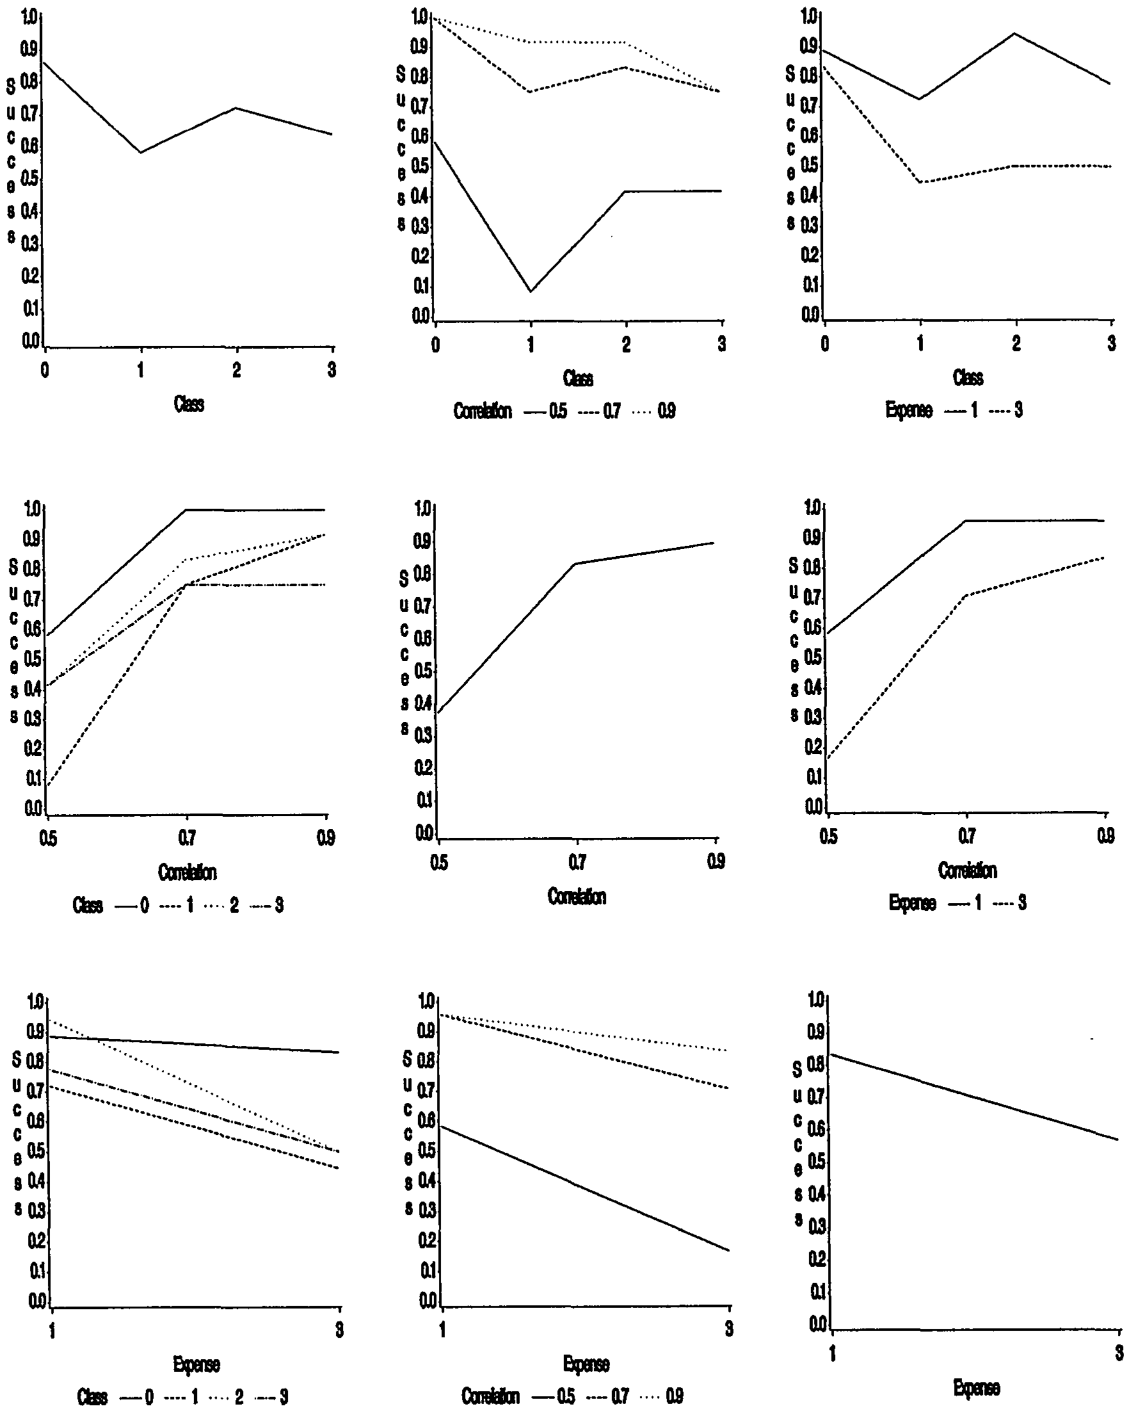

Figure A.1: Success Frequencies for the Look-Ahead Heuristic

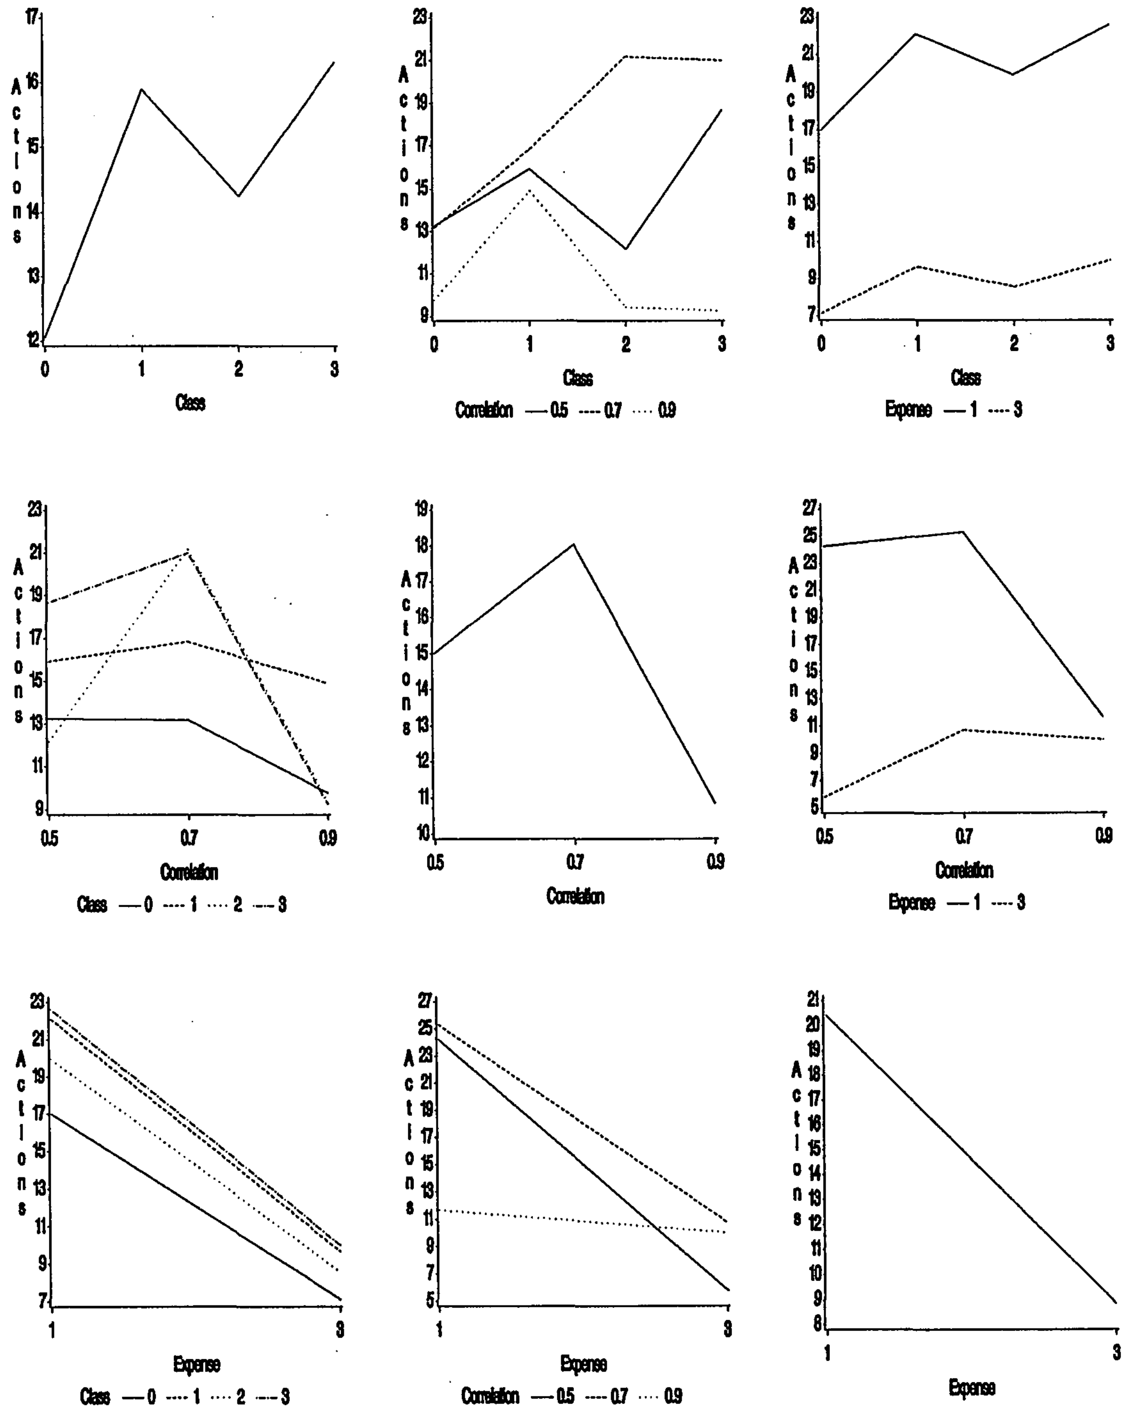

Figure A.2: Experimental Actions for the Look–Ahead Heuristic

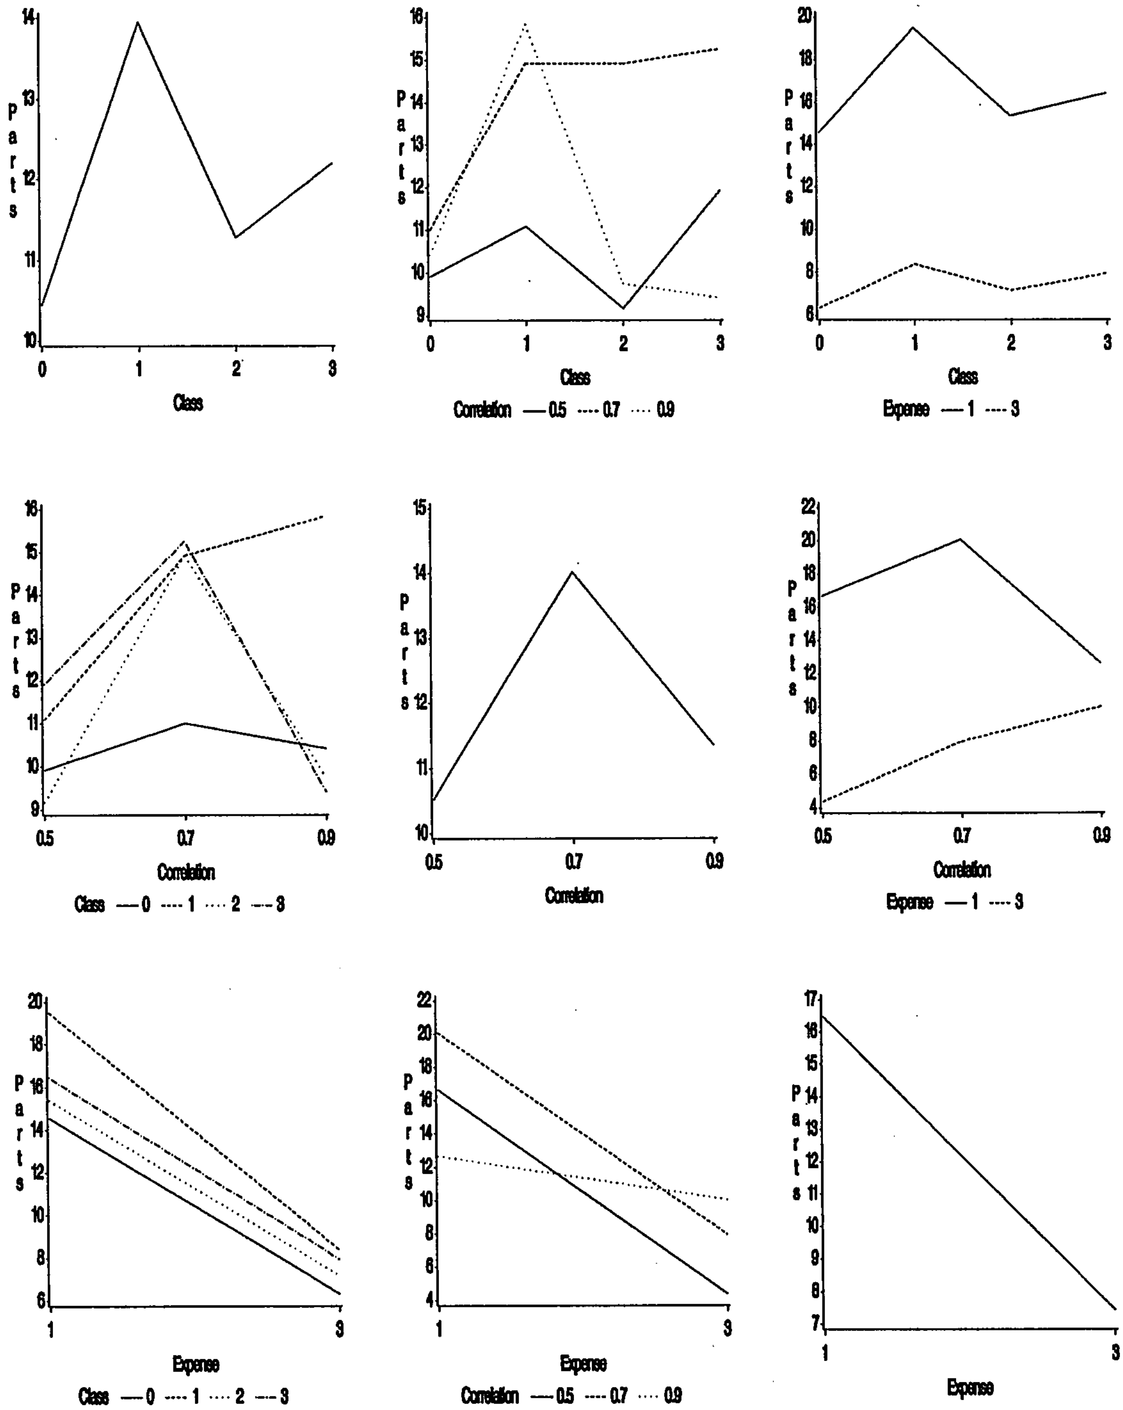

Figure A.3: Parts Used for the Look–Ahead Heuristic

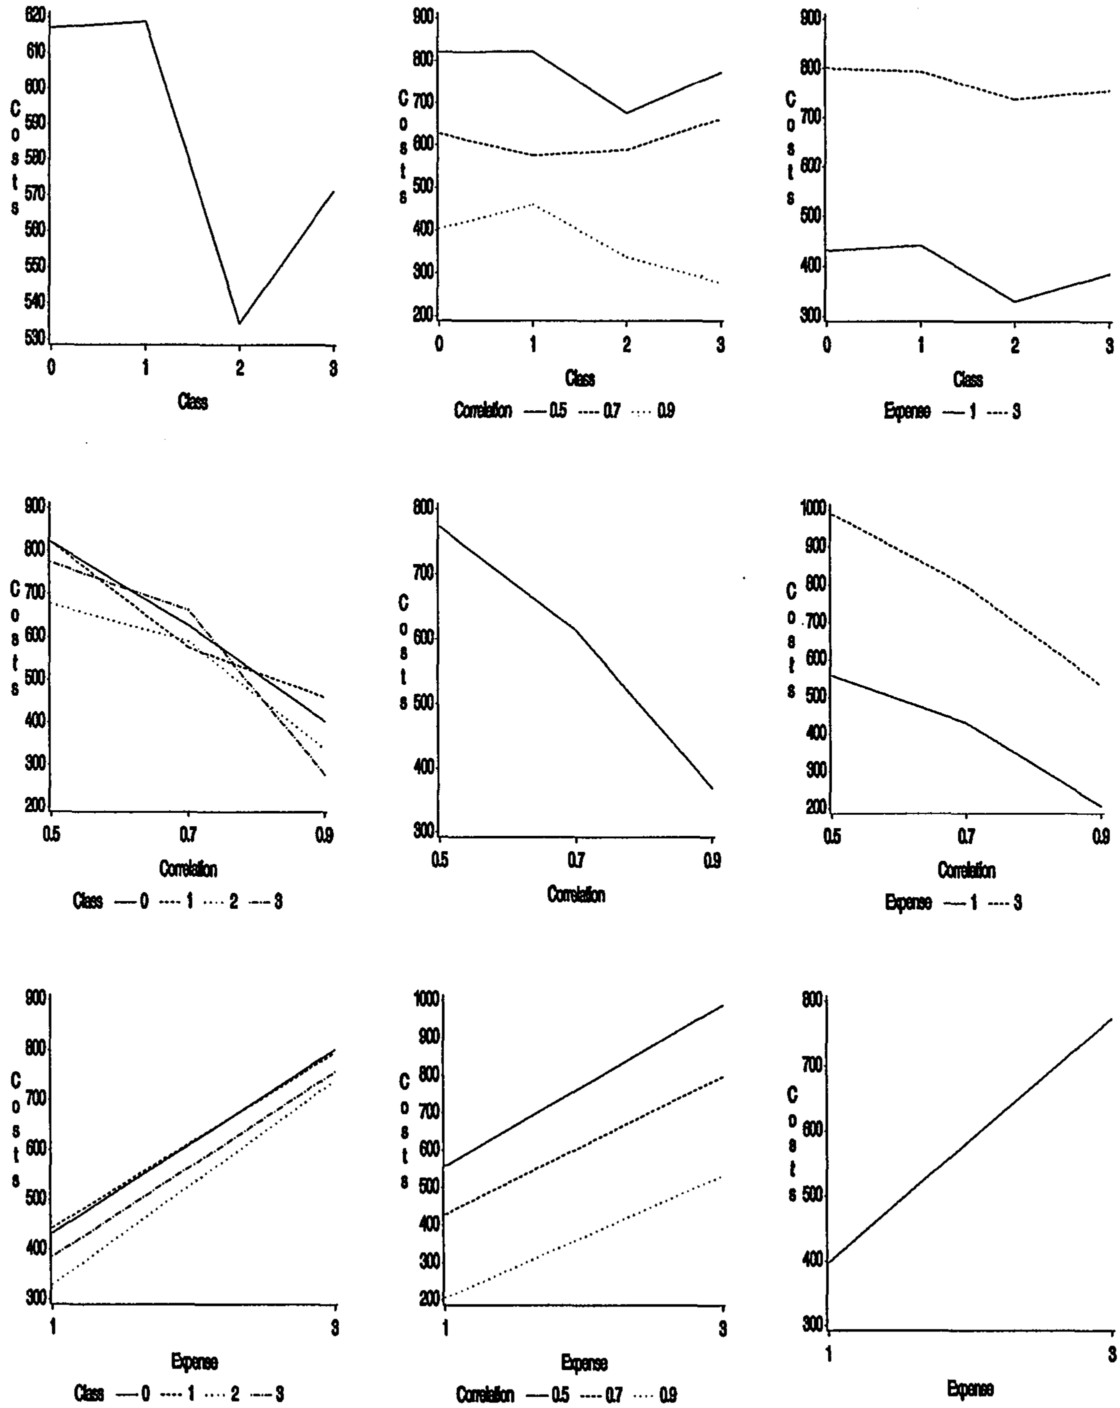

Figure A.4: Total Cost for the Look-Ahead Heuristic

APPENDIX B. RESULTS FOR THE SWAPPING HEURISTIC

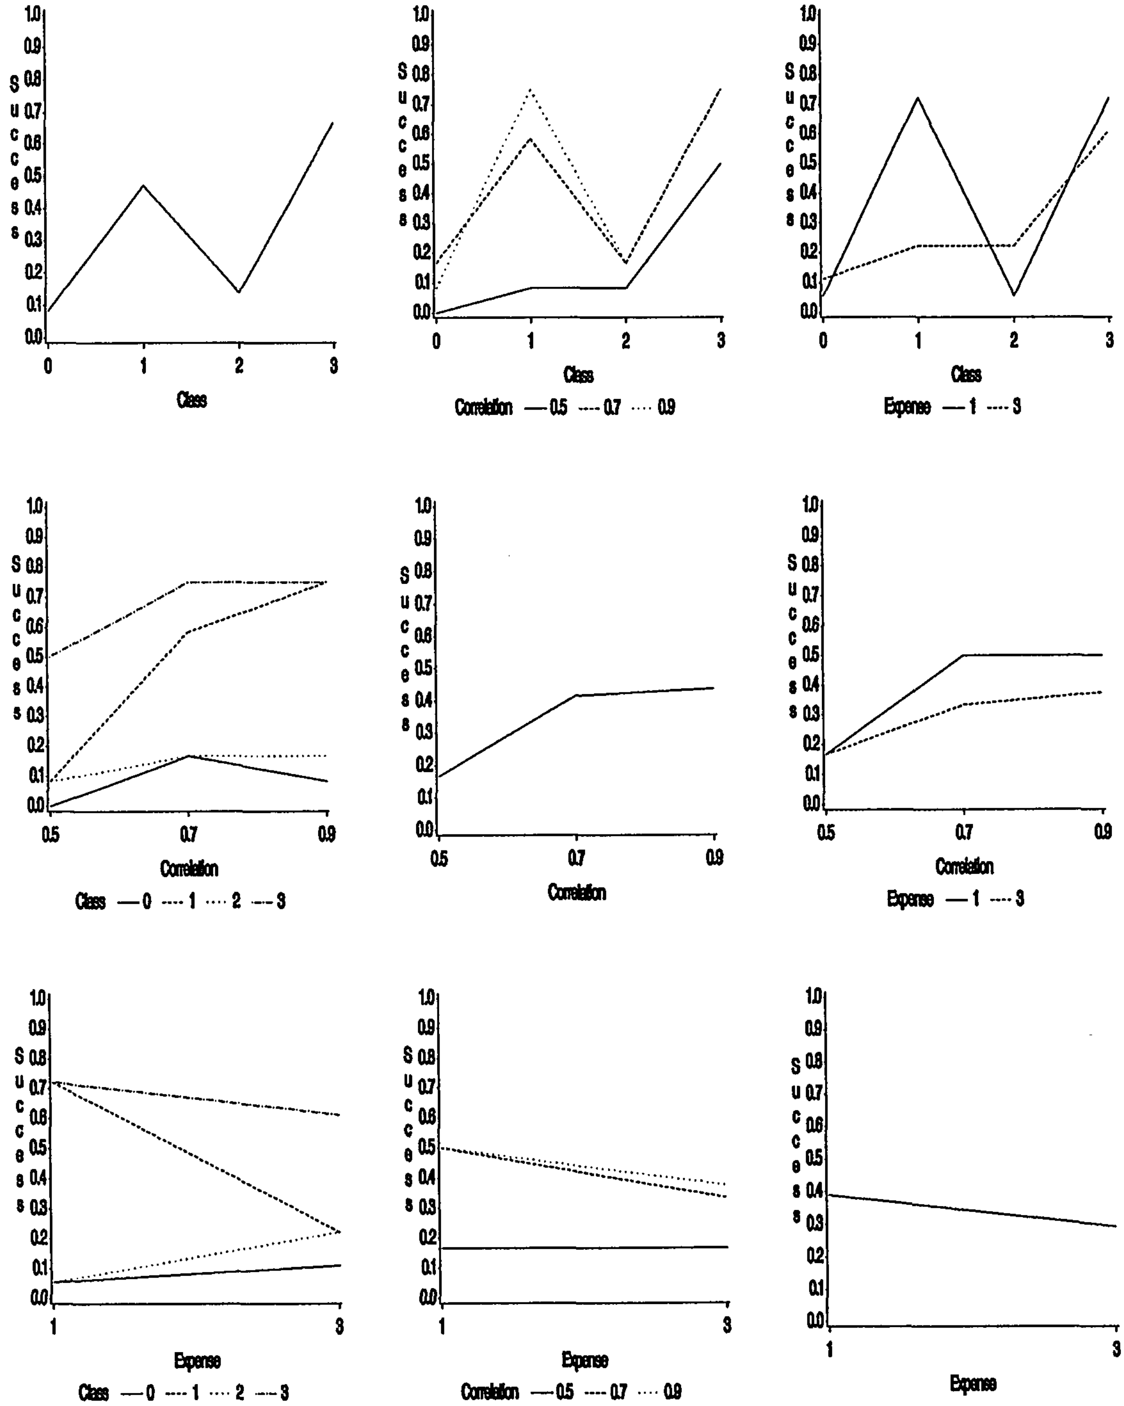

Figure B.1: Success Frequencies for the Swapping Heuristic

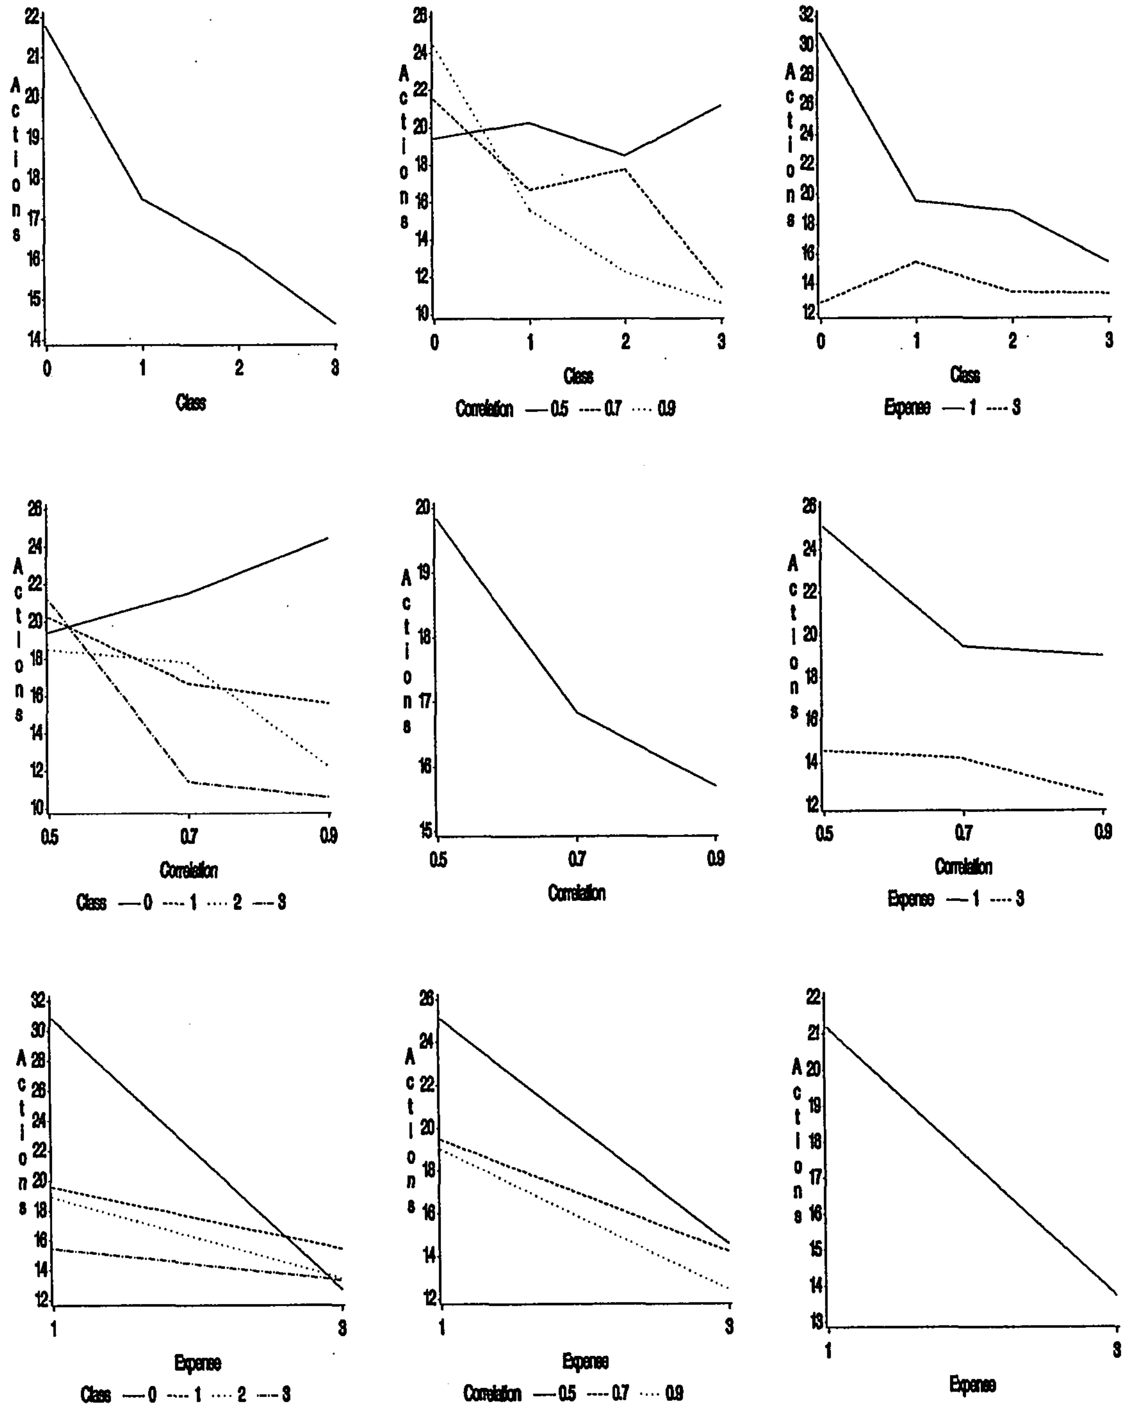

Figure B.2: Experimental Actions for the Swapping Heuristic

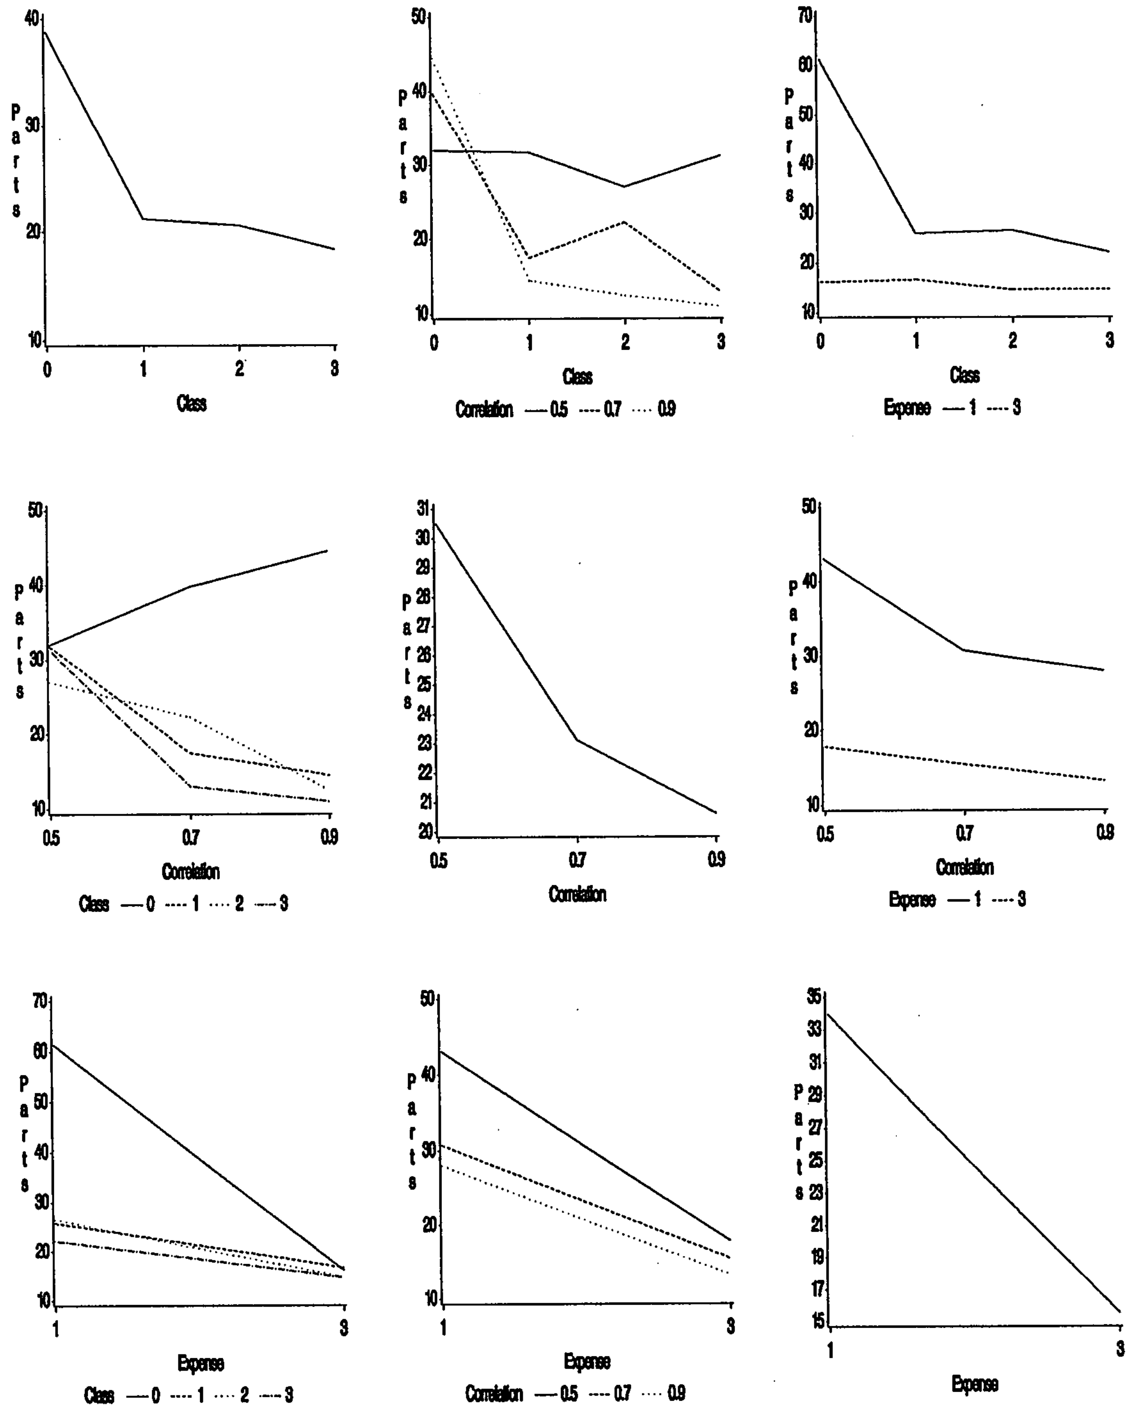

Figure B.3: Parts Used for the Swapping Heuristic

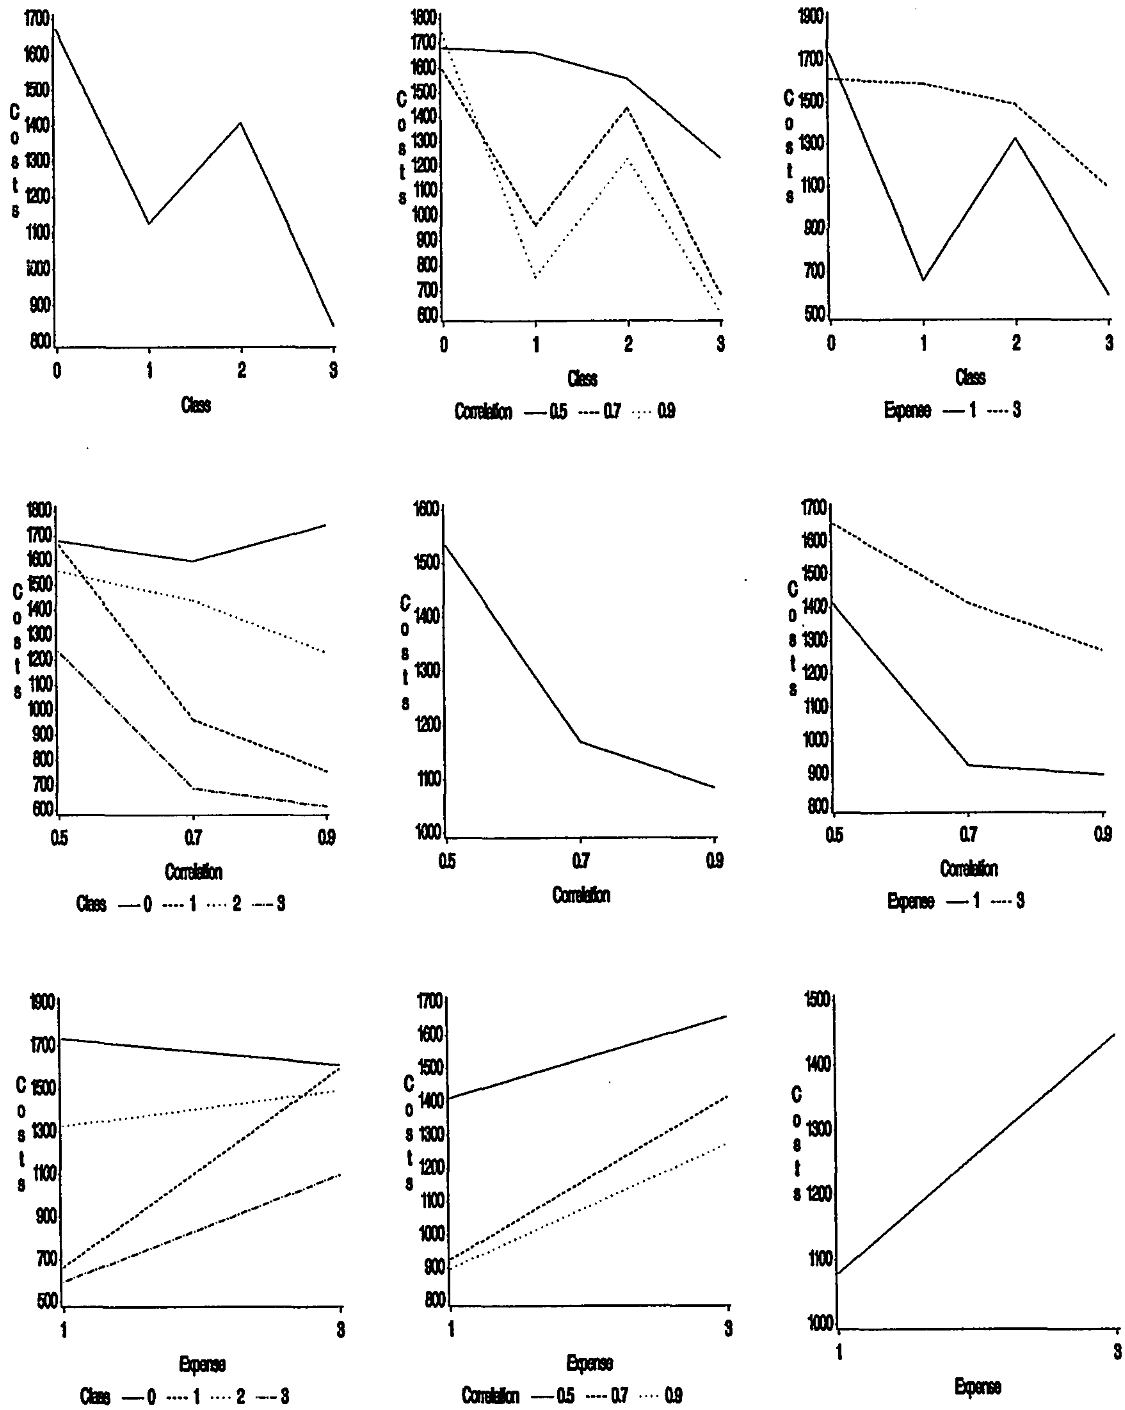

Figure B.4: Total Cost for the Swapping Heuristic